UNIVERSITY OF SOUTHAMPTON

Commitment Models and Concurrent Bilateral Negotiation Strategies in Dynamic Service Markets

by

Ilja Ponka

A thesis submitted in partial fulfillment for the degree of Doctor of Philosophy

in the

Faculty of Engineering, Science and Mathematics
School of Electronics and Computer Science

March 2009

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF ENGINEERING, SCIENCE AND MATHEMATICS
SCHOOL OF ELECTRONICS AND COMPUTER SCIENCE

Doctor of Philosophy

by Ilja Ponka

Technologies such as Web Services, the Semantic Web and the Grid may give rise to new electronic service markets, which have so many services, providers and consumers that software agents need to be employed to search and configure the services for their users. To this end, in this thesis, we investigate bilateral negotiations between agents of service providers and consumers in such markets. Our main interest is in decommitment policies or rules that govern reneging on a commitment, which are essential for operating effectively in such dynamic settings.

The work is divided into two main parts. In part I (chapters 3-8), we investigate how the decommitment policies, through the parties' behaviour, affect the combined utility of all market participants. As a tool, we use decisions that parties make during their interaction. These decisions have previously been discussed in the law and economics literature, but this is the first empirical investigation of them in a dynamic service market setting. We also consider settings (for example, with incomplete information) that have not been addressed before. In particular, we take four of these decisions — performance, reliance, contract and selection — and consider them one by one in a variety of settings. We create a number of novel decommitment policies and increase the understanding of these decisions in electronic markets.

In part II (chapters 9-11), we consider a buyer agent that engages in multiple negotiations with different seller agents concurrently and consider how decommitment policies should affect its behaviour. Specifically, we develop a detailed adaptive model for concurrent bilateral negotiation by extending the existing work in several directions. In particular, we pay special attention to choosing the opponents to negotiate with and choosing the number of negotiations to have concurrently, but we also address questions such as bilateral negotiation tactics and interconnected negotiations on different services. These ideas are again evaluated empirically.

Contents

Acknowledgements

I thank my supervisor professor Nicholas R. Jennings for his supervision and infinite patience, Engineering and Physical Sciences Research Council (EPSRC) for funding this research (grant GR/S62710/01, project 'Virtual Organisations for e-Science') and people at the lab for the help in small practical and technical matters over the years, including Johnsen Kho and Dr. Sebastian Stein for helping me to submit this thesis from abroad.

Nomenclature

β	Beta parameter for time-dependent strategy
χ	Probability that the buyer has a guess about the seller's tactic
γ	Probability that a guess the buyer has about the seller's tactic is correct
$\lambda(t)$	Parameter of the *Poisson* distribution, $i\frac{t_{lastEntry}-t}{t_{lastEntry}}$
ρ_s	Reliability (the probability that the provider s will perform)
a_b, a_s	Probability that the seller (s) or the buyer (b) experiences an adverse effect
$\theta, \theta_b, \theta_s$	Negotiation tactic used by a buyer b or a seller s
$C_s(q_s, t)$	Sellers' cost function
c_s	Cost the seller s has to pay to provide the service
$D(t)$	Probability that the seller has paid its cost at time t
$EC(q)$	Estimated cost function
$EV(q)$	Estimated value function
e	Reliance Level
$F_{min}(q), F_{mid}(q), F_{max}(q)$	Probability that the quality q is below q_b^{min} (min), above q_b^{max} (max) or between them (mid).
f, f_b, f_s	Decommitment fee, for buyer (b) or for seller (s)
i	Seller entry intensity
L_b, L_s, l	The utility decrease caused by an adverse effect
p	Contract price
q	Quality
q_s	Quality of provider s
q_b^{max}	Maximum useful quality for the consumer b
q_b^{min}	Minimum useful quality for the consumer b
r, r_s, r_b	The reservation price, of the provider s or consumer b.
t	Time (turn)
t_0	Beginning of the experiment ($t = 0$)
$t_{c,s}$	Time the seller s has to pay the cost c to provide the service

$t_{e,s}, t_{e,b}$	Time the seller s or buyer b entered the market
$t_{x,b}$	The deadline for the consumer b
$t_{contract}$	Time the contract was entered into
$t_{decommit}$	Decommitment time
$t_{delivery}$	Delivery time ($t = 1000$)
$t_{lastEntry}$	The last turn the seller's are allowed to enter the market
t_{min}, t_{max}	The start and end time (increasing decommitment policies)
$U_b(q, p)$	Utility of a single buyer (consumer)
$U_s(p, c_s)$	Utility of a single seller (provider)
U_{b+s}	Utility of a single buyer and seller (parties to a same contract)
$V_b(q)$	Buyers' value function

Chapter 1

Introduction

The Internet has grown very big and keeps growing at a very fast pace. No human can possibly keep up with this much information and if the information is not in machine-readable format, much of the potentially relevant information may be impossible to find with machines too. This is one of the problems the Semantic Web (Berners-Lee et al. 2001) is seen to address. Thus, when information is encoded in machine-readable format, it becomes possible for machines to do more sophisticated searches and find better, more relevant information for us humans to use. Similar ideas are also used in the area of Web Services (Curbera et al. 2002), which can encapsulate the implementation details (operating system, programming language and so on) and allow existing services to be offered to the other computer systems, even to outsiders.

On the other hand, nowadays, many types of science, and increasingly also businesses need to analyse large amounts of data. Computers make it easy to collect vast amounts of information, but using that information will often need sophisticated analysis tools and when these tools are applied on such vast amounts of data, large quantities of computing power are also required. Although computers have become both cheaper and more powerful in the last few decades, the price of the most powerful computers (nowadays often a cluster of 'normal' computers) remains high and maintaining (and constantly upgrading) these powerful systems requires specialised know-how and is usually quite expensive. For many companies and universities the returns for such investment would either be uncertain or too small to justify the cost. Thus, Grid technologies (Foster and Kesselman 2004) aim to alleviate this problem by allowing parties to 'rent' computing power and specialised tools and procure all sorts of services from other parties online when they need them. These services can range from simple infrastructure services such as computing power, storage capacity and communication bandwidth, to complicated service packages which involve aggregating many types of information and simpler services from multiple sources into one comprehensive service. In addition, these services can allow customisation to each customer's needs to a varying degree.[1]

These technologies have the potential to create huge electronic service markets. Such markets will be very dynamic and complex environments. With possibly thousands of services, many of them configurable in many ways, with hundreds of different providers, and thousands of service consumers,[2] from all over the globe, the market will contain so much information and so many possibilities and change so rapidly (the service providers and consumers coming and going), that no mere human will be able to make the necessary informed decisions quickly enough to fully use their potential. Therefore, in these new markets, it will be machines, not humans, interacting with other machines. The humans will still decide the goals of these interactions, but it will be the machines that take care of the details automatically and mostly behind the scenes.

Since these services will require resources to produce and will be useful to their users, it is reasonably certain that the providers of these services will be seeking a renumeration for their efforts.[3] On the other hand, since at least some of these services will be configurable, the parties will have to negotiate on the details of the service before they are actually provisioned.

Given all this, the best way to describe such markets is as a multiagent system, where each human or company is represented by their own agent (an encapsulated computer system that is situated in some environment and is capable of flexible and autonomous action in this environment in order to meet its design objectives (Jennings 2001)). These agents then negotiate with each other on behalf of their masters.

Given this vision, in this work we will concentrate on such negotiations. Obviously there is a huge number of questions and issues with such negotiations and we can only hope to make progress in some of them. In the following, we will discuss our approach and interests and why we chose them over some others. However, we

[1] Mass customisation (as introduced by Pine (1993)) has been seen as one of the potential key benefits of electronic commerce (for example in Vulkan (2003)).

[2] Or even billions in some visions, see for example Kephart et al. (2000).

[3] Much of the World Wide Web works now with advertising income and is free to the end-users, but in these markets that strategy is unlikely to work, since there are no guarantees that their masters would ever see any advertisements that were given to them.

start this introduction by explaining the background for our work in more detail (section 1.1). This will be then followed by the aims for our research (section 1.2), a summary of our contributions (section 1.3) and the outline for the rest of the thesis (section 1.4).

1.1 Implementing Dynamic Service Markets

We start by discussing some of the key characteristics of the marketplaces we are interested in in this work. In particular, the key properties are:

- *player autonomy*: All entities in the market are autonomous, in the sense that they have their own agendas and interests. This means that there is no one centre of control, rather different entities make their own decisions and the marketplace is an institution where conflicting interests are settled.
- *player heterogeneity*: The different providers offer different services, have different costs, resources and situations and different consumers need different things from their services, and have different goals and plans.
- *openness*: The population of market players is not limited in advance, but any party that is technically able to interact with the marketplace and its parties and that is willing to accept the basic rules of the marketplace, is allowed to enter. New market players can enter and existing ones can exit the market at any time without any warning.
- *dynamism*: The circumstances of the market players can change at any time and that, in turn, may change their goals and strategy. For example, a person looking for a two-week trip abroad for a holiday, but a sudden problem with a boiler in his house means that he suddenly has less financial resources to spend on a holiday and he switches the plan to spend a weekend in a neighbouring city instead.

These characteristics make it difficult to use traditional approaches to software development (Huhns and Singh 2005). So, a number of new approaches and technologies are being developed to operate in such circumstances. We discuss some of these approaches in section 1.1.1. We then discuss the need for negotiations between the agents (section 1.1.2) in these markets. We conclude by discussing the central role of commitments (and decommitments) in section 1.1.3.

1.1.1 Background

We start by introducing two key technologies for our approach. We first discuss Service-Oriented Computing (SOC) (section 1.1.1.1), which is a popular framework for describing market situations such as ours. We then go on to describe agent-based computing as a way of adding autonomy and flexibility to the traditional view of service-oriented computing (section 1.1.1.2).

1.1.1.1 Service-Oriented Computing

In Service-Oriented Computing, each entity (*service provider*) offers one or more *services* to those who need them (*consumers*). Each such service comes with a description which explains what the service will do, how it should be used and what policies it will follow. In more detail:

- *Expected effect*: An important characteristic of the effect descriptions is that they describe only *what* the service does, but not *how* (Huhns and Singh 2005).[4] From the consumer's point of view, a service is therefore opaque, since the implementation details are hidden. The idea is that the consumer does not need these to decide if he wants to use the service or not.

- *Interaction*: The interaction description contains all the information needed for using the service properly in a standardised, machine-readable language (Huhns and Singh 2005). Also the interface itself uses some implementation neutral technology (such as Web Services (Curbera et al. 2002)), which allows the provider to implement the service with any programming language, in any operating system, using any libraries, databases and other tools he wants.[5]

- *Policies*: The policies explain restrictions, constraints and requirements for the use of the service (Huhns and Singh 2005). These cover issues such as how the service consumer must be identified, how his authority to use the service will be verified and how the communications between the parties will be encrypted.[6]

[4]Huhns and Singh call this property coarse granularity and according to them its main advantage is that it reduces dependencies among participating services.

[5]Huhns and Singh call this property implementation neutrality.

[6]Many organisations and contexts have implicit policies that govern their processes. Huhns and Singh see the writing of explicit policies as a useful exercise that makes it easier to implement and enforce company-wide policies. The explicit policies that contain the relevant rules also make it possible for automated inter-organisational processes to take place.

The main design point of service-oriented computing is therefore *interoperability*. It should be easy for any consumer to use any service in the market and it should be easy to combine different services into larger service packages. In particular, the consumers can use the services from the market as a part of their internal processes or they could combine the services from the market and/or their internal processes into service packages that they can sell in the market. Another important goal of service-oriented computing is *scalability*. Since the services are produced in the provider's own information systems and since most of the communication occurs directly between the parties, the architecture should be able to scale, even to all of the Internet. The only potential bottleneck is the market itself and especially the service repositories, from which the the consumers can find the service descriptions and search for the services they need.[7] However, there can be multiple separate repositories and the big repositories can be replicated, so these should not prove too problematic either.

However, arguably the most interesting characteristic of service-oriented computing is that of *dynamic service selection*. This means that due to all the interoperability characteristics, the consumer can select the service it wants to use at run-time. Thus, instead of binding the program to always call the same service (or a small group of services), the programmers can write their software to dynamically select the service from the service market, whenever the need arises. This makes it possible to write more flexible and fault-tolerant programs. If one service fails to respond, the program can automatically locate other similar ones and use them instead.[8]

1.1.1.2 Agent-Based Computing

It has been argued that agent-based systems are well-suited for managing large and complex distributed systems (Jennings 2001) such as these service markets. Here an *agent* is viewed as an encapsulated computer system that is situated in some environment and is capable of flexible and autonomous action in this environment in order to meet its design objectives (Jennings 2001).

[7] The service repository is not a required element, but since there must be some way through which the consumers can learn of the services, it is usually present.

[8] It also decreases the risk of a lock-in, i.e. a service provider exploiting the consumer's dependency on his service, for example by increasing the price of using it. If the binding is dynamic and automatic, the program would, in such circumstances, just switch to another provider.

In more detail, the service-oriented computing paradigm we just described can be applied to agents directly or the service-oriented architecture implemented by other technologies (such as Web Services, for example) could be complemented by agent technologies. The *direct application* would mean that we describe our service market as a multi-agent system, where each service provider and service consumer has one or more agents representing their interests. The services would be provided by agents to other agents.[9] The *complementary approach* would mean that the services themselves are not agents, but for example Web Services, and the role of agents is to help their masters to find the right services and possibly combine them to useful packages (Huhns 2002). In both cases, the agents can also negotiate between themselves on the price, quality and any number of other properties of the service.[10]

In this work, we are interested mainly in the negotiation between the provider and consumer agents, which happens before the actual service provisioning. Therefore this service provisioning is, to us, of secondary importance. Moreover, since negotiation can be easily incorporated into either of the two models discussed above, we do not need to choose between them. Thus, it does not matter in the negotiation phase, if the service is going to be provided by the provider agent or through some other means (such as a Web Service).

In summary, our service markets basically consist of a large number of autonomous actors, each with their own goals and plans. This is likely to make the market a very complex and dynamic environment for the agents.[11] Specifically, there will be agents entering and exiting the market, there will be agents using different tactics and strategies, having different parameters and/or agents needing different services. Moreover, since the agents' owners are human, or companies run by humans, their goals and plans can change at any time. This means that the instructions they give to their agents can, in theory, change at any time without any advance warning and this can radically change how the agents behave. In

[9] Although the discussion of service-oriented computing often assumes the use of Web Services and other related technologies, they are by no means a requirement. Any technology that can provide the properties we described above could be considered to be service-oriented.

[10] In the complementary approach this would mean that each provider would have a set of negotiating agents that would negotiate with the consumer agents. The right to access the service could then be made contingent on the acceptance of one of the provider's negotiating agents.

[11] The division into dynamic and static environments is common in the agent literature (for example Russell and Norvig (2003) and Wooldridge (2002)). A static environment is one that can be assumed to remain unchanged except by the agent's actions. In contrast, in a dynamic environment there are other processes operating on it and hence the environment changes in ways that are beyond the agent's control.

this uncertain context, each consumer agent must be able to navigate and find the services its owner needs, for reasonable prices, often within a specified time (i.e. before a deadline). Thus, the interactions between the agents are central to this vision and it is there that we explore next.

1.1.2 Interactions in the Market

As already mentioned, the services that are traded in the market are going to be useful to their consumers and, on the other hand, there may be costs associated to their production. Given that the market players are independent, they are likely to be selfish.[12] A selfish provider will try to sell its services for profit. It is therefore highly likely that the services are exchanged for some renumeration. Here, we will assume that the renumeration will always be money.[13] Given this, the parties must have some way to agree on the terms of exchange. We discuss here the three main options, namely fixed and dynamic pricing (section 1.1.2.1), auctions (section 1.1.2.2) and bilateral negotiation (section 1.1.2.3).

1.1.2.1 Dynamic Pricing

The price of a service in dynamic service markets can be set very much like in many real-life markets of today: the seller posts the price (one price for everybody) and the buyers either take it or leave it. However, the electronic environment allows for much more dynamic price-setting than that. It is much easier (cheaper) for the seller to change the prices from time to time or even customer to customer. Both types of variation have of course been used for a long time, but an electronic environment brings new possibilities in both. Because in electronic environments, the price is nothing more than a value in a database (or something equivalent), it can easily and at very low cost be changed whenever necessary. In contrast, in traditional brick-and-mortar shops, the owner might need to print a new menu or catalogue, change price tags or reprogram cash registers and that is both very

[12]In the real world, which is dominated by separate agents (humans and companies), altruistic behaviour is quite rare, especially in the commercial contexts. Even ‘free’ goods and services are almost always paid for by some one else (for example an advertiser) or the receiver of the ‘gift’ himself later in higher fees (for example in UK, a consumer signing for 12-18 month contract with a phone operator and getting ‘free’ phones in the process will obviously pay for the phone in the higher monthly installments).

[13]We do not consider barter (exchanging services for other services), because that would introduce an unnecessary complication and would not bring anything essentially new to our problem. For similar reasons, we are not interested in how the money actually changes hands or how services are actually delivered.

time-consuming and expensive. The lower costs may mean smaller and more rapid price changes. Moreover, in an electronic marketplace, the small pricing changes can also be left (at least partially) to a computer program that can take into account things like stock situation and demand in real-time (if this information is available).

However, electronic environments also allow more data to be collected from each and every customer. In a small shop of the old times, the shopkeeper knew all his customers and their needs and preferences, sometimes better than the customer themselves. However, in the huge supermarkets of today, this is simply not practical. People shop in more than one place, in many shops self-service is becoming more and more common (both in product selection and in paying) and they may meet a different clerk at the counter every time.[14] Now, computer systems can track each and every customer even in this sort of environment. A supermarket chain might use a loyal customer bonus card or something similar to connect certain purchases to a certain customer and also get information about when they did their shopping and how they paid for it. This information can then be used in many different ways. In an electronic marketplace, however, the customer can be tracked even before he buys anything, through the products he wants to see, searches he makes and everything he has ever bought from this store. And this information could, in theory, be used to make profiles out of customers and even setting prices to a level that this particular customer (and ones like him) would find acceptable. Naturally, this could be a different amount for different types of customers.

So the electronic service markets offer the sellers a possibility to collect information about the buyers and make it easier (cheaper) to experiment with the pricing. However, this does not mean that the seller is able to set an optimal price in every situation and with every customer.

1.1.2.2 Auctions

The possibilities of price-setting go beyond this type of one-sided scheme. Instead of trying to guess what the customer would be willing to pay for a given service (given all the information available), a service provider could just ask them. One simple way for doing that is to organise an auction. For example, Google uses

[14] And the clerks are too busy and meet too many people to keep track of who is buying what.

auctions to set prices for the advertisements on its search results page. The advertisers offer Google money to show their ads in searches that have certain words in them and the highest bidders get the service (i.e. Google puts the highest bidders on its result page). The price is always what the advertisers are willing to pay.

Auctions do have some desirable properties (such as efficiency, neutrality, formality and simplicity that make them possible to analyse), but they also have some drawbacks (such as inflexibility and difficulties in controlling information) and, in addition, there are some situations, in which they simply do not work properly. Indeed, many important results of auction theory (see Krishna (2002) for an introduction) are only valid under very narrow assumptions. We discuss these and other questions in more detail in section 2.3.1.4. At this point, it is sufficient to note that there is a real need also for other approaches.

1.1.2.3 Sequential and Concurrent Bilateral Negotiations

In this work, we concentrate mostly on bilateral negotiation. In a market setting, where more than one negotiation is possible, the multiple negotiations can occur either one after the other (sequentially) or at the same time (concurrently).[15] We will investigate both. On one hand, we investigate many-to-many markets where interaction occurs through repeated one-on-one negotiations (sequential bilateral negotiation) and, on the other hand, we investigate concurrent bilateral negotiation, which means that each agent takes part in multiple one-on-one negotiations at the same time for broadly the same service with a goal of entering into just one contract. In the latter setting, we will concentrate on one-to-many negotiations because it is the base case.

The main advantage of negotiations is that they offer a possibility for *two-sided interaction* during the interaction. In such interactions, both parties can indicate what they want and what they are willing to give in return. This will alleviate the problem of incomplete information and can lead to better outcomes for both parties. The provider can get a better understanding of his customers' needs and use this information to make his selection more interesting to his potential customers.[16] The consumer gets a service that matches his needs. This effect is especially strong in multi-attribute situations, where a service has many attributes that can be configured interactively. Such configuration clearly requires two-way

[15] Of course combinations of these two are also possible.

[16] Through negotiations the provider may also get information on why his offerings are *not* competitive. This information may be quite difficult or expensive to obtain by other means.

interaction between the parties, especially if there are too many possible configurations to list all the possibilities as separate services. However, also in pure price negotiations, the parties can learn what the other party is expecting and willing to pay for a certain service.

Another setting in which bilateral negotiations could prove to be useful, is interconnected negotiations. That is, situations where the agent is after multiple services and these services have connections to each other. For example, the consumer might either want two different services or none at all (complements) or he might want one of the three services (substitutes). Now, such dependencies can be managed in specific type of auctions, but these can become too complicated to manage if the number of services and dependencies is high. In contrast, if interconnectedness is managed with concurrent negotiations, the agent and its master can flexibly manage the complexity of the task themselves. We will discuss these issues more closely later, in section 2.3.1.4.

Negotiations are also much more flexible (changes in the environment can quickly be taken into account), they allow rich and flexible information exchange (in cases where the service required or options available are not clear for the participants) and they offer some strategic advantages (e.g. the possibility of differentiating strategies between different opponents and of changing information between different negotiations). The strategic advantages are more apparent when negotiations are conducted concurrently. On the other hand, concurrent negotiation typically requires more computation and communication resources than a simple auction. We will discuss these issues in more detail in section 2.3.1.

At this time, negotiation, especially concurrent bilateral encounters, is probably not crucial to many contemporary e-commerce systems. However, it is easy to see that in these service markets it is likely to become so. Since the services are configurable and different configurations are likely to need different amounts of resources and to be of different values to the consumer, some form of negotiation is clearly required to establish the appropriate details of the service, and the renumeration to be paid. The configurability would suggest negotiating on the multiple attributes (different characteristics of the service plus price) and it is true that this would be a setting especially suitable to concurrent bilateral negotiation, since this type of interaction is difficult to organise in an auction.

On the other hand, the ability to reach agreements through negotiation is a fundamental capability of intelligent autonomous agents (Wooldridge 2002; Lomuscio et al. 2003). When we consider also the amount of information and time constraints involved, it is clear that only a machine, a software agent, can accomplish the task and therefore automated negotiation conducted by software agents will be the means of accessing the services these markets offer.

1.1.3 The Role of Commitments and Decommitments

In a multi-agent system there are dozens, sometimes even hundreds of software agents, each of which may have its own goals. In a competitive setting, like an electronic market for services, these agents compete against each other for the providers' services or the custom of the consumers. Contracts (or commitments about performing a service at a later time) are an often used tool in such systems to coordinate the agents' behaviour and to enable individual agents to make assumptions about the future actions of other agents. In other words, the contracts provide a degree of predictability to counteract the uncertainty caused by the distribution of control, and each agent having its own goals and decision-making facilities (Jennings 1993). On the other hand, in the field of real-life commerce, enforcable commitments make the exchange of goods or services and money safer in situations in which the parties perform their part of the bargain at different times. For example, if provision of a certain service would require time and resources from the provider before it can be delivered and if there are no commitments, the provider would have to start working on the service and hope that the consumer is willing to pay for the service when he finishes. The provider can of course ask for the price to be paid in advance, but, without commitments, it would then be the consumer who would have to trust that he will actually get something in return for his money.

In fast-moving electronic markets, software agents can form hundreds of contracts in a matter of seconds (He et al. 2003). On the other hand, in dynamic, ever-changing environments, sometimes the parties of these contracts may live to regret their choices. A better alternative may surface only seconds after agreeing to something or the circumstances may change and turn a contract from lucrative into disastrous. The full commitment contracts (that allow no reneging) can be very problematic in such systems and force parties to perform contracts that have become very onerous (because the circumstances have unexpectedly changed). In some cases, it is also possible that performing according to the contract is simply not possible anymore, in which case an unexpected failure may occur. So, the problem is that although commitments in general are useful and necessary, there are cases where they may be counter-productive. We will first discuss the main approach to this problem in agent-based markets, leveled commitment contracts, and shortly their alternatives (section 1.1.3.1). However, the problem of setting an optimal decommitment fee (amount to be paid by the decommitter to the victim) remains open and we hope that more detailed analysis in the law and economics literature will be able to give us some pointers (section 1.1.3.2). Much of this work is about investigating these ideas in dynamic service markets.

1.1.3.1 Leveled Commitment Contracts

Different approaches to tackle the commitment problem have been suggested. In *contingency contracts* (Raiffa 1987), for example, the contract itself would (ex ante) specify the circumstances under which the commitment is no longer binding. In practise, it may be very difficult to list all possible situations in which this would be necessary and it is difficult to draw the line in cases where the situation can become progressively worse. Another approach used in real-world international commerce is *renegotiation* (Craswell 1988), which means that after the change in circumstances (ex post), the parties renegotiate the contract to take the change into account. Unlike contingency contracts, renegotiation allows parties to react to unforeseen contigencies. The problem here is that it relies on the other party's cooperation: A party can always reject an offer to renegotiate and is under no obligation to accept any changes to the contract. Both of these approaches share a problem, which is that it may be difficult to reliably establish that a cause for re-negotiation or dismissal of the contract actually exists and one party is not just claiming so to force the other party to accept worse terms. This can be a problem especially if the change has to do with one party's internal circumstances. In multi-agent systems, a third option is often used, namely Sandholm and Lesser's (1996) *leveled commitment contract*, which:

- allows unilateral decommitting for both parties at any time, but
- requires the decommitting party to pay the opponent a monetary fee (called a decommitment fee) for doing so.

In this mode, the reason for decommitment is not relevant, but any reason will do and this can be kept private. There is no need to arrange proof that any change has occurred. This simply allows a party to abandon a contract that has become counter-productive; that is, its utility has become negative. Now, because the party must pay the penalty in order to decommit, a decommitment occurs only if the decommitment improves the decommitter's utility more than the fee to be paid reduces it.

1.1.3.2 Performance, Contract, Reliance and Selection Decisions

In the literature that discusses leveled-commitment contracts, the decommitment fee is mostly seen as a deterrent of decommitment. This view, however, overlooks the effect a decommitment has on the other party (the victim) and on the society in general. In particular, the decommitment will usually decrease the utility of the victim, because he will lose the profit he was expecting. Moreover, it is possible that he has already accrued some costs (preparing for its performance) before decommitment occurs. When the contract is abandoned, these efforts may become useless. Now, if these lost profits and accrued costs outweigh the benefit the decommitter receives from decommitment, the decommitment actually decreases the sum of utilities of the parties and is, therefore, detrimental to the welfare of the society as a whole.

In contrast, the law has traditionally taken another view, that of the victims. In cases of non-performance, the victim (in most legal systems) is entitled to damages[17] and the aim of damages is usually to put the victim financially in the same position as if the contract had been performed appropriately (Treitel 2003). That is, the damages compensate the victim for the loss that the non-performance causes. The economic efficiency of this rule has been investigated in the law and economics literature in the area of efficient breach theory (Barton 1972). The conclusion of this work is that this is the optimal policy from the society's point of view. In particular, by setting the damages (decommitment fee) equal to the damage caused by the decommitment to the other party, a breach (decommitment) occurs when and only when the benefit to the decommitter is greater than the damage to the victim. Therefore decommitment always increases the total welfare of the society.

However, the economic analysis does not stop there. In the law and economics literature, the role the damages (decommitment fees) play in parties' decision-making (and through those decisions on the welfare of the society) is seen as more complicated than that. Different authors classify these decisions slightly differently

[17]There may be other remedies in some situations and in some legal systems, but we concentrate on damages here.

and identify different numbers of them,[18] but the fact remains that the damages do affect the optimal behaviour of the parties in many ways. In this work, we will investigate the following four decisions (Smith 2004; Craswell 2001; Kornhauser 1986):

- *Performance*: Whether or not to perform an existing contract. The higher the level of commitment (the higher the decommitment fee), the more likely the parties are to perform. Too high or low levels of commitment can be counterproductive for the common good (as explained above).
- *Reliance*: How much to rely on the performance of the other contract party. The less likely the other party is going to perform, the less the performance should be relied on. Relying on a performance will often cause costs and the more the performance is relied upon the higher the costs may be. In case of non-performance, these costs may be wasted.
- *Contract*: Whether or not to enter a contract at all. Although a contract may be beneficial now, the circumstances may change later and make a contract counterproductive. The higher the required level of commitment (the higher the decommitment fee) and the higher the probability of such adverse changes in circumstances occuring, the less inclined a rational agent should be to enter a contract in the first place. From the society's point of view, very risky, high cost contracts may be inadvisable.
- *Selection*: Who to transact with. For example, an opponent's reliability should be taken into account when choosing transaction partners. Here, non-performance does not produce any benefit to the society, but only a performance does. Any costs invested in a service that was never delivered, are wasted. If the other party's level of commitment (the decommitment fee) is too high, the party might be indifferent between performance and non-performance or it may even prefer non-performance (so either be indifferent between different reliabilities or even prefer less reliable providers).

The problem with these different decisions is that the optimal amount of damages (decommitment fee) they prescribe are different. As explained, in the terms of performance decision, the optimal fee would fully compensate for the loss of

[18] Craswell (2001) identifies no less than seven different decisions and says that even his list is not exhaustive but that any decision taken by the parties in relation to the transaction is, to some extent, affected by the extent of the liability they and their opponent bear.

the opponent. On the other hand, in the terms of selection decision, such a fee would make a party indifferent between performance and non-performance and therefore the party could potentially choose very unreliable providers (or at least not the optimal providers) that would be likely to lead to a suboptimal result for the society, even overall losses. Of course if the opponent in this case knows that he is not likely to perform and the decommitment fee is high, he might not enter a contract in the first place (contract decision). All these decisions are therefore very intertwined and the only way to find an optimal decommitment fee for a given market would be to take into account all the effects the fee has taken into account. In many real-life markets this might prove to be very difficult, even impossible, given the complexities of such systems and the incomplete information they operate with. Given this, any solution would be highly environment-specific and, in any non-trivial environment, finding an optimal policy might not be possible analytically and the only way of doing it would be through experimentation. However, one should be aware of these effects when allowing decommitments and using decommitment fees. In this work, our aim is to investigate these decisions and how they affect market behaviour.

There are also two views to each of these decisions. On the one hand, each and every one of them can be seen as a means of optimising the common good. The society will be better off if all beneficial contracts are performed, appropriate reliance is put on the performances and very high cost contracts that have very small chance of actually ever happening are avoided. On the other hand, they can also seen as a guideline for an individual agent for making decisions in dynamic service markets. These decisions will tell the agent designers what factors he needs to consider when his agent is making decisions. First, an agent will only perform its contractual duties, if it is still in its own interest (benefit from performing is greater (or loss smaller) than from not performing and paying a decommitment fee). Second, the amount an agent relies on the performance of the contract partner depends on how likely is the performance, what is the benefit from performance and what is the benefit/loss from non-performance. Third, whether or not the agent enters in a certain contract at a certain time depends on whether the expected benefit is greater from entering or not entering that contract. And finally who to negotiate and transact with depends on who of the potential contracting partners offer the best expected utility, given the value of the service they offer, price likely to be required to secure the service and reliability of the provider.

As can be seen, the first view is about maximising the common good and the other one is about maximising the individual good. As we said, there are many decisions and they all are interrelated. So are these two views. To be able to maximise the common good, one needs to know how individual agents behave and what factors they will consider. The economic analysis of damages (fees) usually assumes that the agents take decisions that maximise their expected utility (view 2) and fees are set so that this leads to maximisation of common good (view 1). However, in real-life systems it may simply be too complicated for this approach to work and it may be impossible to find an *optimal* decommitment policy for a given setting.[19] On the other hand, it might well be possible for a single agent to make its own decisions in a near-optimal manner (depending on the information it has available). We therefore use both views. On one hand, we investigate the effect of different decommitment policies and the parties adapting to them on the common good. On the other hand, we also investigate how consumer agents in a more complicated setting (concurrent bilateral negotiation) can adapt to their environment, even if finding an optimal policy might not be possible. We now turn to discuss our approach in more detail.

1.2 Research Aims

In this work, we investigate the effect that commitments and decommitments, especially decommitment policies (the marketplace rules about decommitments), have on the common good and on individual agents' strategies in a dynamic service market, in which the individual circumstances of the agents (or their masters) can change at any time.

To this end, we will have two different settings with clearly separate themes to investigate:[20]

- *Markets with Subsequent Bilateral Negotiation*: In this setting, we will have a simple marketplace, in which many buyers and sellers of a certain service meet to exchange services for money. Every buyer is matched to a single seller at random. The pairs negotiate on the price of the service for a while. This process is repeated a few times. The buyer and seller agents will be as simple as possible and the main interest will be how decommitment policies affect market-level characteristics, especially the total utility of all parties (the common good). The view here is usually that of a benevolent system

[19] If direct experimentation is possible, a reasonably good policy may well be possible.

[20] These two settings mirror the two views on performance, reliance, selection and contract decisions discussed in the previous subsection.

designer, one who is implementing the marketplace and trying to maximise the common good of the participants.

- *Concurrent Bilateral Negotiation*: In contrast, in concurrent bilateral negotiation, we will concentrate on one consumer agent and its strategies in a dynamic and complex market with many sellers. We give our agent the power to choose the opponents it negotiates with and ask it to choose the number of opponents it will negotiate with concurrently and conduct the negotiations while taking into account many variables including the decommitment policy and the probabilities of circumstance change. We will also shortly discuss strategies to deal with concurrent negotiations on different, but interconnected, services. The view here is that of a designer of an individual agent.

Given this context, we will now discuss the key requirements (R1–R6) for our settings which differentiate our approach from the others. All of them have to do with different types of decisions that the parties make in relation to a transaction.

R1. Costly and Time-Consuming Service Preparation
In order to provide a service to a consumer at a certain time, the provider must invest resources and time in its preparation before the delivery. In more detail, the seller has to start preparations for the delivery at time $t > 0$ before the actual delivery and he has to pay the cost c_s when he starts the preparations. Once paid, this cost is not recoverable. Such a cost means that the society and the sellers can be adversely affected by decommitment (i.e. be in a worse situation than he would have been had he not entered into the contract in the first place). They may have incurred costs that may turn out to be wasted when the consumer decommits. This has connections, for example, to the contract and performance decisions.

R2. Changing Circumstances
The circumstances of the parties to a contract can change after the contract has been entered into. The change can occur to one or both parties at a random time between the contract and delivery time. The change is always adverse and decreases the utility of a contract to the party by amount a. Such a change can transform a lucrative contract into a disastrous one. This requirement is connected to the contract decision.

R3. Heterogeneity
The providers and consumers can be heterogeneous, meaning that they have

individual characteristics such as quality they produce (or need) and these can vary from one provider (or consumer) to the next. The utility the parties to a contract obtain depends on these characteristics and therefore for the providers or consumers, exactly the same contract can provide a different utility. This requirement helps to make the market slightly more difficult to analyse and predict, but it also makes considering selection decision a necessity where there is a large disparity between buyer and seller populations and not everybody can find a transaction partner.

R4. Limited Negotiation Resources

Negotiation, especially concurrent negotiation, can be resource-intensive. This means a consumer agent in concurrent bilateral negotiation settings to limit the number of negotiations in which they engage.

R5. Openness

New providers and consumers can enter the market and old providers and consumers can exit at any time.

R6. Availability of Limited Basic Information

The marketplace has reliable information about all the participants in the marketplace and it will give this information to all participants for free. We assume for example that the quality of service and reliability of the provider are (accurately) known by all consumers. This requirement is made because we do not plan to investigate how the parties collect this information or deal with incomplete information in this area. Our task is to understand how decommitment policies affect welfare of the society and optimal negotiation strategies in many different ways and this problem, although interesting, would not help us in addressing that question. On the other hand, we assume that some information remains private (for example negotiation deadlines) and some is not know by anyone before the related event happens (for example, if and when circumstances will change for a given participant).

Obviously we expect our agents to be capable of adapting to these variations that the first five requirements offer. As explained, the purpose of the sixth requirement is to restrict certain research problems from this work.

1.3 Research Contributions

In this section, we summarise our contributions to the state of the art. As per the rest of this thesis, it is divided into two different parts: the market setting (with subsequent negotiation) and the concurrent negotiation setting.

In the market setting, we will investigate the effect the different decommitment policies have on the common good. We will offer a new and more complete view on the role, problems and possibilities of decommitment in dynamic service markets. Specifically, we will investigate each of the four decisions we discussed earlier (performance, reliance, contract and selection) in a dynamic service market setting and see how parties will adapt their behaviour to consider these decisions and then investigate empirically how this adaptation affects the common good in a market with many buyers and sellers. Although the work in this part builds on existing models and principles from law and law and economics, these models and principles have never before been applied empirically to a dynamic service market setting with subsequent bilateral negotiations, where the parties may need to decommit, or consider incomplete information or sub-optimal policies. In the process of investigating settings with incomplete information and the relative performance of sub-optimal policies (limits of the optimal policy's performance superiority), a number of novel decommitment policies are developed and investigated. In particular,

C1. with the performance decision, we will consider the effect of one and two-sided decommitments, re-entries and incomplete information. This involves discussing each setting in detail and developing new decommitment policies for each of them, often looking at legal rules for inspiration. Here, usually compensating for the actual loss of the victim is the key.

C2. with the reliance decision, we consider the boundaries of compensating the victim's loss in a setting where the buyer (victim) can decide how much to rely on the performance. Based on earlier literature, we show in several different settings that compensating for extra-reliance will lead to overreliance and, subsequently, a loss in common good. We contrast the buyer taking the reliance decision into account with several simple reliance strategies in different settings. We also investigate several novel decommitment policies that offer partial restrictions on compensation (inspired by law, again) but find that they are usually not effective compared to the case where the decommitment policy offers no compensation for extra reliance.

C3. with the contract decision, we investigate especially the possibility that the parties can use the contract price as a risk-allocation tool between themselves, allowing them to find mutually acceptable deals while ensuring non-negative expected utility for them under any circumstances and under any decommitment policy. We also investigate the effect of none, one or both parties taking possible changes into account and the limits of this mechanism when decommitment policies, settings and available information vary.

C4. with the selection decision, we investigate how the decommitment policy affects the buyers' choice between different negotiation and contract partners, which have a varying level of reliance as well as quality, and how these choices affect the common good. We investigate several near-optimal policies and compare their performance with the optimal and other policies.

In the concurrent bilateral negotiation setting, we consider a case where we have one buyer agent and several seller agents and the buyer agent can be negotiating with many sellers at the same time. Such models have existed in the literature, but we improve and extend them by:

C5. making the design of the concurrent bilateral negotiation model such that each different question is isolated to a separate and interchangable module and defining the interaction of these modules in sufficient detail that the operation of the whole is clear in all circumstances. The modular structure allows new strategies and tactics to be easily added.

C6. allowing the sellers to use several different bilateral negotiation tactics, some of which make it necessary for the buyer to adapt its behaviour to them (even under complete information) or at least consider that it will be successful in some negotiations and not so successful in others. The buyer's tactics designed to counter these (and other similar tactics) will be novel and aimed at one or more of the possible opponent tactics and they will take into account the possibility of decommitments (the contract decision). We will investigate both cases with full information about opponent tactics and cases with limited or unreliable information.

C7. considering, for the first time, the essential issues of opponent selection (the selection decision) and concurrency control (choosing the number of opponents to negotiate concurrently) in concurrent negotiation. We also make it clear how these two interact with each other, the negotiation strategy (choosing what tactic to use in each negotiation) and other parts of the model. Moreover, we discuss situations where the buyer agent knows as well as situations where it does not know the negotiation tactics the seller agents are using.

C8. considering simple interrelated negotations on different services simultaneously in cases where the different services are substitutes (only one needed) or complements (all of them are needed). All interconnections between services consist of different combinations of these basic cases we investigate.

C9. using empirical data (instead of known distributions) on previous runs to estimate what sort of offers to expect later in the experiment and whether or not to accept the offer on the table or wait for a better one later.

In the end, we will have an adaptive concurrent bilateral negotiation model which will be able to adapt to more variation than anything in the state of the art and we will have investigated several new tactics and strategies empirically in several different settings.

1.4 Thesis Outline

After this introduction, we will first review the relevant literature in relation to our requirements (*Literature Review (Chapter 2)*). We identify the models and parts of models that we can use in our work. In addition, we highlight some shortcomings of current research that we need to address in our work. After this, the thesis is divided into two major parts each consisting of several chapters.

In part I (chapters 3–7), we will discuss commitment models. Our focus is on common good and we try to find the effect different decommitment policies have on the combined utility of the market participants. We will discover that the effect of these policies is very significant. Specifically, we will describe our model for a dynamic service market and then investigate the effect each of the four decisions (performance, reliance, contract and selection) will have on the common good of the market participants in turn. In more detail:

- *The Marketplace Model (Chapter 3)*: We explain our basic model for a dynamic service market and explain the basic decommitment policies.

- *The Performance Decision (Chapter 4)*: We discuss situations where the utility of one or both parties may be adversely affected after the contract has been entered into and the party in question has to decide whether to perform or decommit and the effect the decommitment policies have on that decision. We discuss settings with complete and with incomplete information and with and without a possibility of re-entry.

- *The Reliance Decision (Chapter 5)*: We analyse situations where the buyer has to decide how much to rely on the seller's performance in a given contract, when the seller's reliability (probability of performance) is known and the role the decommitment policies have on that decision. We discuss settings with different reliability distributions.

- *The Contract Decision (Chapter 6)*: We detail how the parties can use the contract price as a means of risk distribution in different settings and how they can, by considering whether or not to enter into a contract, safeguard their own interest in all circumstances and settings.

- *The Selection Decision (Chapter 7)*: We discuss the buyer's decision to choose its negotiation and contract partners and how that decision is influenced by the decommitment policies. We will again have different seller reliability settings.

In part II (chapters 8–11), we discuss concurrent bilateral negotiation strategies. We will develop a three-layered model for adaptive, interconnected, concurrent bilateral negotiations and then discuss each layer in more detail. In other words:

- *An Adaptive Concurrent Bilateral Negotiation Model (Chapter 8)*: We will first introduce our model and the environment it needs to work on. The model is modified from the marketplace model (chapter 3).

- *Negotiation Tactics: the `Negotiator` Level (Chapter 9)*: The `Negotiator`s manage single bilateral negotiations. We discuss different negotiation tactics they face and different tactics they can use to tackle them.

- *Managing Concurrency: the `Controller` Level (Chapter 10)*: The `Controller`s manage a set of `Negotiator`s negotiating on a single service. Here, managing concurrency refers to choosing who to negotiate with, how many negotiations to have in a given moment in a given setting and what tactics to employ in each negotiation. We introduce several basic strategies for managing these problems.

- *Coordinating Concurrent Negotiations: the* ***`Coordinator`*** *Level (Chapter 11)*: the `Coordinator` manages a group of `Controller`s that negotiate on different, but interconnected, services. We discuss the two basic cases of interconnection: subsitututes (where only one service is needed) and complements (where all services are needed) and discuss some basic strategies.

Finally, in *Conclusions and Future Work (chapter 12)*, we will offer our conclusions and explain some possible directions for future research.

Chapter 2

Literature Review

In this chapter we discuss the relevant literature and pinpoint the state of the art in the relevant questions. Since we cannot discuss all the literature relating to (automated) negotiation (simply because there is too much of it), we limit our discussion to the literature that is especially relevant to our approach. This review consists of three main parts. First, we discuss decision and game theory, which are present throughout our work (section 2.1). We then proceed to discuss the literature that is relevant to each of our two settings in turn. We discuss literature for the market setting and decommitments in section 2.2 and the literature for the concurrent bilateral negotiation setting in section 2.3. We then summarise our findings in section 2.4.

2.1 Decision and Game Theory

We start this literature review by investigating decision and game theory, which will be extensively used in this work. In more detail, *decision theory* offers mathematical tools for a single decision-maker to make decisions under uncertainty, when he is the only one making the decisions or the other decision-makers can be abstracted away from the situation. When there is more than one autonomous, often self-interested, party interacting usually *game theory* (Fudenberg and Tirole 1991) is used instead. Game theory is a branch of microeconomics and can be used to analyse situations and find optimal strategies for all parties involved. Usually the objective of game theoretic analysis is to find an equilibrium where all parties have their set strategies that they do not want to change. We will make heavy use of both decision and game theory and, therefore, we start our literature review by reviewing some basic concepts.

The starting point for *utility theory* is a set of possible (mutually exclusive) alternatives from which the individual has to choose (Mas-Colell et al. 1995). The second requirement is that the decision-maker must be able to rank the outcomes that follow from his choice. In more detail, there are two properties:

- *completeness*: given any two outcomes, the decision-maker must be able to say whether he prefers one or the other or if the outcomes are indifferent to him.
- *transitivity*: if the decision-maker prefers A to B and B to C, then he must prefer A to C.

Now, a *utility function* assigns a real number (*utility*) to every possible outcome such that more preferred outcomes get assigned larger numbers than less preferred ones. It is worth mentioning that the game-theoretic utility function is ordinal in the sense that it will rank all the outcomes (A > B > C), but its value (utility) does not give any information on how much better A is to B.[1] Strictly speaking, this means that summing the utilities of different decision-makers is a meaningless exercise and, therefore, the sum of utilities cannot generally be used as a measure of the common good or the welfare of everybody in the market or society. Instead, a concept of *pareto-optimality* is often used: a solution is pareto optimal, if nobody's utility can be increased without decreasing somebody else's utility. However, in automated negotiation, the sum of the agents' benefits or utilities is sometimes used as a measure of global optimality, usually as a secondary criterion among pareto optimal solutions (Rosenschein and Zlotkin 1994). Since the parties are selfish (not interested in the social optimum), this measure can only be used at design time. In addition, the utility functions of the different players have to be cardinal and therefore give some idea how much better one deal is than another.

In this work, we are mostly interested in price negotiations between two parties. Fortunately, this is one situation where there is a simple way to create utility functions that can be considered to be cardinal. For the provider, the utility function in the literature is often of the form:

$$U_{provider}(p_c, q) = p_c - c(q_s),$$

[1]Any function that ranks the outcomes in the correct order is a valid utility function.

where p_c is the final price and $c(q_s)$ is the cost of producing the service of quality q_s. For the consumer, the utility function is:

$$U_{consumer}(p_c, q) = v(q_s) - p_c,$$

where $v(q_s)$ is the value of the service with quality q_s (or the maximum price the consumer is willing to pay for that service) and p_c is the price paid for it. Quite often, the value function $v(q)$ is very simple, $v(q) = q$. The utility function then operates as a measure of success for the different players. The higher the utility, the more the agent has succeeded in the negotiation. It can also be argued that these utility functions are cardinal.[2]

Now, we started this discussion on utility by saying that it will help us to make decisions under uncertainty. So far we have not discussed decision-making. However, using the utility functions to make decisions is straight-forward. In a simple case, where each action leads to one outcome, the decision-maker just selects the action that leads to the most preferred outcome (i.e. the outcome with the highest utility). However, the power of this approach becomes evident when we consider situations where there are multiple possible actions, each of which can lead to many possible outcomes with known probabilities. We can extend this simple model by replacing the utility function with an *expected utility function*, which assigns all the actions an *expected utility*, which for action a is simply:

$$EU(a) = \sum_{k=1}^{n} z_{a,k} U(o_k),$$

where $U(o_k)$ is the utility of outcome o_k, $z_{a,k}$ is the probability of outcome k occurring given that action a was taken, and n is the number of possible outcomes. After the expected utility is calculated for all possible actions, the decision-maker selects the action that has the highest expected utility. This rather simple tool (maximising the expected utility) is a very powerful method for making decisions under uncertainty. The obvious problem with it is that it requires that the probabilities of different outcomes given any of the possible actions are known. However, we assumed that we will have information of this type (requirement **R6**), so this approach will be used at least in some form.

[2] For example, let there be two providers with the cost of 0.1. One of them sells his service for 0.3 and the other for 0.5. Hence the first one makes a profit of 0.2 and the second one a profit of 0.4. It could be argued that the second one has done twice as well as the first one. Similar arguments could be made in the case of consumers. Indeed, Rosenschein and Zlotkin (1994) use similar logic in their utility functions and consider them to be cardinal.

2.2 Commitment Models in Dynamic Service Markets

In the first part of our work, we will be interested in the common good and how decommitment policies can affect it. To investigate this, we will have a marketplace where many buyers and sellers meet, are matched at random and negotiate bilaterally on the price of a service. The mechanisms of the marketplace are not very important, it is the decisions about entering into contracts, performing contracts and all related decisions that matter. However, we have to choose some mechanics and we will therefore shortly discuss automated bilateral negotiation, especially the three main approaches to it (section 2.2.1). However, our main interest in this part is the notion of commitments and decommitments: how they affect the common good in the market and through what mechanisms. We start this discussion by considering changes in circumstances and three different approaches (contingency contracts, renegotiation and leveled-commitment contracts) that can be used to manage commitments when the commitment can become very onerous or difficult to fulfil (section 2.2.2). Having found the leveled-commitment contracts to be the most promising option for us but having found no reasonable guidance on how to set the level of commitment well in different settings, we investigate some decisions that parties make during their interaction to see if they would be able to provide us with a useful starting point (section 2.2.3).

2.2.1 Bilateral Negotiation: The Main Approaches

We start by discussing automated *bilateral negotiation*, which has been a popular research topic over the years and there is a lot of literature on it. We do not plan to discuss all of the literature here, but we concentrate on the three basic approaches to automated bilateral negotiation, namely game theoretic, heuristic and argumentation-based. This trichotomy was introduced by Jennings et al. (2001) and although it is not conclusive (there are some other approaches), these three present a clear majority of all papers on autonomous negotiation. We introduce each and discuss their applicability to our problem in turn. We start with the game-theoretic approach (section 2.2.1.1) and follow it by the heuristic (section 2.2.1.2) and the argumentation-based approaches (section 2.2.1.3).

2.2.1.1 The Game-Theoretic Approach

The *game theoretic approach* uses game theory (as discussed in the previous section) to analyse the negotiation situation and to find the optimal strategies for both negotiators and preferrably an equilibrium where no party wants to change their strategy. The game theory literature on negotiation is very rich and wide (see Muthoo (2002) and Osborne and Rubinstein (1990) for an introduction). However, much of the work builds on Rubinstein's analysis (1982) of the *alternating offers protocol*. In this protocol one party makes the first offer for the contract and after that the recipient of an offer has always three possible actions, it can either:

- accept the offer,
- make a counter-offer or
- quit.

The basic limitation of the model is that after making an offer, a party cannot make a new offer (or quit or accept), until it has received a counter-offer from the opponent.[3] Negotiation ends with a contract, if any offer is accepted by the opponent and without a contract if one of the parties quit.

In theory, parties could continue making counter-offers forever. However, there are many possible ways to introduce pressure to end the negotiation one way or the other in finite time. In game-theoretic models, the typical method is *discounting*: the value of the object of negotiation decreases over time, so the longer the negotiation, the less there is to divide between the parties.[4] Another method that is sometimes used in economics is *fixed negotiation costs*. Such costs force the parties to compare the chance of getting a better deal and the cost of continuing the negotiation. At some point, the cost will be greater than the expected benefit and a party will accept whatever is on the table (or walk away).

In the automated negotiation literature, two other sources have also been used: deadlines and competition. A *deadline* is the time by which the agreement must be reached, if it is to be reached at all. Clearly both agents will try to end these

[3] The model does not include the possibility of cancelling an offer, so a party is bound to his offer until it has been rejected (the other party has quit or made a counteroffer). The other party can therefore create a contract simply by accepting an offer. We will not consider offer cancellation in this work.

[4] Discounting means that the value of the negotiation object is multiplied by a discount factor, which is a real number between 0 and 1 (typically close to 1) after every turn of negotiation.

negotiations successfully before the deadline. The deadline is usually different for different players and the parties usually do not know each others' deadlines. This causes uncertainty, since the opponent's deadline can be now and his current offer may be his final one. If, for example, party A makes a counter-offer, his opponent will either accept it, make a counter-counter-offer or withdraw (if it was his final offer). A will have to balance these possible outcomes with his other option, which is to accept whatever is on the table. As time goes on, the probability that the opponent withdraws usually increases while the offer on the table improves. This means that the temptation to accept what is on the table increases.

Competition can also have a big impact on the strategies. For example, if there are many providers, but only a few consumers, the providers must impress the consumers quickly enough, so that the consumers will not go elsewhere. This will mean that the providers will make reasonable offers very quickly in an effort to ensure a deal. If the services were identical and information was complete, this would lead to full price competition, which would drive the prices to marginal costs.[5] However, neither of these assumptions hold in our environment. Our services are heterogeneous (requirement **R3**) and the parties' information about the market is incomplete (requirement **R6**).[6] Therefore, some providers would be able to make a profit even in such circumstances, although obviously the providers would be in difficult position.

Now, Rubinstein's seminal paper provides unique optimal strategies for the parties in an alternating-offers negotiation with an infinite number of rounds, discounting and complete information. When both players play optimally the negotiation always ends in the first round when one party makes an optimal offer, and the other party accepts it. However, this classic model is not directly applicable to our problem, since it assumes complete information and the negotiation ends in acceptance of the first offer. Both are problematic. The agents in our problem do not have complete information (requirement **R6**) and the optimal offers in the first round would not leave much room for negotiation or later changes (requirement **R2**).

[5] This is the result of the basic price competition model from economics (Varian 2003). With different marginal costs there is a twist, however. If there were k consumers, the k providers with the lowest marginal costs would set their prices a bit under the costs of the provider with the $(k+1)$th lowest cost.

[6] The providers do not know each others' costs and the consumers do not know all prices in the market due to their limited negotiation resources (requirement **R4**). Spulber (1995) has shown that the former is enough to prevent full price competition. On the other hand, Janssen and Rasmusen (2002) showed that already a non-zero chance of being the only provider the consumer knows of, will have the same effect. In our setting, some consumers might only negotiate with one provider, so a chance that a provider would be the only one is positive.

Fortunately, it seems that these two problems are linked. In the later literature, when two-sided incomplete information about the opponent's reservation price has been introduced into the model, negotiations do not usually end in the first round. Unfortunately, building and especially analysing these models is a very difficult task (Cho 1990) and often some unrealistic assumptions have to be made in order to make the analysis possible at all. It is common, for example, to allow only one of the sides to make offers (the other can only accept or reject, see for example Cho (1990), Cramton (1984) and Ausubel and Deneckere (1992)), to allow parties to withhold their offer (use the delay as a signal of private information, for example Admati and Perry (1987) and Cramton (1992)), to allow only two different reservation prices (Chatterjee and Samuelson 1987) or to limit the length of the negotiation to two rounds (Fudenberg and Tirole 1983). Almost all the models use discounting as a method of putting pressure on parties and one that uses a fixed cost for negotiation (Perry 1986) ends up with the negotiation proceeding at most one round. And in spite of these restrictions, many of these models have multiple possible outcomes, in which case a sensible outcome must be identified (Cho 1990).

In addition, the game theoretic models are not very robust. Even the smallest change in the environment may require recalculation of the strategies, which can be quite a demanding task. Moreover, our environment is very dynamic, so the only constant property of the environment is the change. Therefore, a purely game-theoretic approach to bilateral negotiation does not seem realistic in our case. Therefore, we will not use the game-theoretic approach as a basis for our bilateral negotiation tactics. We will, however, apply game and decision theory in many other ways.

2.2.1.2 The Heuristic Approach

Agent environments in general are often quite complicated, but still the agents need be able to make decisions that improve their chances of achieving their goals. Now, a very typical way to solve this problem in these situations is to use a *heuristic approach* (Jennings et al. 2001). Such approaches acknowledge that it is usually impossible to find an *optimal* solution in such a complex environment, so instead the goal is try find a *good enough* solution.[7] In negotiations, a heuristic approach

[7]According to Simon (1955), this sort of heuristic behaviour is actually very typical for human decision making. A more recent overview of this problem is given in Aumann (1997) and one approach on these 'satificing' games is offered in Stirling (2003).

means that the negotiation space (possible deals) is searched in a non-exhaustive fashion and that the number or type of different strategies is limited.

In this vein, Jennings et al. (2001) have listed the key advantages and problems of the heuristic approach. The advantages are that the models can be based on more realistic assumptions (such as incomplete information) and also work in more complex situations, where mathematical analysis is impossible. Moreover, the use of heuristics allows designers to use less constrained models of rationality. The problems are that the models tend to generate sub-optimal outcomes and that there is no guarantees of success. In other words, the heuristic approaches require extensive empirical analyses and simulations, because it is usually impossible to predict precisely how the system will behave and the outcomes that will arise in different circumstances.

We will now discuss some of the main heuristic strategies; namely time-dependent, resource-dependent and behaviour-based in more detail. These three were recognised in Faratin et al. (1998).[8] First, in *time-dependent tactics* the predominant factor used to decide which value to offer next is time t. The basic idea is to vary the acceptance value[9] depending on the remaining negotiation time. In all variations, the acceptance value starts at some high value (for the sellers) or low value (for the buyers) and ends equal to the reservation value at the deadline. The expressions for the acceptance values using these tactics is for the seller:

$$x_S(t) = max - \alpha(t)(max - min)$$

and for the buyer:

$$x_B(t) = min + \alpha(t)(max - min).$$

where max is the maximum offer, min is the minimum offer, t the current time, and $\alpha(...)$ a function that gives a real number in $[0, 1]$, zero at first and one at the deadline. In this case, a wide-range of time dependent strategies can be defined by varying α. However, usually the monotonic variations are used. In particular, three monotonic variations are used so often that they have their own names, boulware, conceder and linear. When using the boulware tactics, the agent tries to hold on to a high price, but when the deadline gets close, the offer drops rapidly to the reservation price. In contrast, an agent using conceder tactics concedes quickly

[8] They explicitly say that their grouping is not exhaustive and neither is the specification of the groups themselves. There are other groups and other possible strategies within each group.

[9] An acceptance value is an offer that the agent makes on its next turn. It is also a threshold for accepting any offers made by the other party: any offer that is better than the acceptance value will be accepted.

in the beginning to near its reservation price and then decreases the price slowly. Linear tactics are in between: the agent makes the same concession at every turn. In all of these cases, α is an increasing function and is usually described as either:

- polynomial: $\alpha(t) = \left(\frac{t}{dl}\right)^{\frac{1}{\beta}}$ or
- exponential: $\alpha(t) = e^{\left(1-\frac{t}{dl}\right)^{\beta}\ln\kappa}$,

where t is the current turn and dl the deadline. In addition, both versions have a parameter β that determines how fast concessions are made and when. The exponential version also uses parameter κ, which sets the starting level for the bids.[10] These time-dependent strategies are the most used heuristic tactics in the literature (e.g. Fatima et al. (2004), Li et al. (2005), Mok and Sundarraj (2005) and Hou (2004) to name but a few).

Second, the *resource-dependent tactics* are similar to the time-dependent ones. Indeed, time-dependent tactics can be seen as a type of resource dependent tactic in which the sole resource considered is time. Whereas time vanishes constantly up to its end, other resources may have different patterns of usage. For this family, there are two broad approaches. Either making the value of t_{max} dynamic or making the function α depend on an estimation of the amount of a particular resource.

Third, in the *behaviour-dependent tactics* the agent takes the lead from its opponent: if the opponent concedes, then it will too. They work best in situations in which the agent is not under a great deal of pressure to reach an agreement and in cooperative situations. Three different variations were introduced in Faratin et al. (1998): relative tit-for-tat, random absolute tit-for-tat and averaged tit-for-tat. In *relative tit-for-tat* the agent reproduces, in percentage terms, the behaviour that its opponent performed k rounds ago. If the agent is a seller and the opponent buyer increased its bid by 25 %, the agent will decrease its own bid by the same amount. In *random absolute tit-for-tat*, the agent does the same, but uses the same absolute amount and in addition, a small random term is added. So, a seller's decrease of £2, would be met with a buyer's increase of £1.90-2.10, for example. In *averaged tit-for-tat* the agent calculates the average of percentages of changes in a window of size $j \geq 1$ of its opponent's history. So, for example if $j = 5$ and the

[10] Clearly $\alpha(0) = \kappa$.

seller has decreased its price from £2 to £1 in five turns, the buyer will make a 10% increase for its offer.[11]

As discussed above, in part I we are not really that interested in the details of the bilateral negotiation process itself. For us, it is sufficient that this process produces potential and actual contracts and that factors like changing circumstances (requirement **R2**) and heterogeneous opponents (requirement **R3**) can be easily taken into account. And heuristic negotiation tactics are simple to use and understand, so we use them in our work as a basis for our bilateral negotiation. As also discussed above the time-dependent negotiation tactics are possibly the most used negotiation tactic in the literature and it can provide us with the flexibility we are after, so we use it in part I on both sides.[12]

2.2.1.3 The Argumentation-Based Approach

Both game-theoretic and heuristic approaches have two limitations (Jennings et al. 2001). First, the proposals are only single points in the negotiation space. Second, the only feedback that the agents get from their proposals is a counter-proposal, which itself is another point in space, an acceptance or withdrawal. This may be very problematic; especially in multi-issue negotiations, where the space of possible contracts is huge.

Now, the *argumentation-based approach to negotiation* aims to remove these limitations (Rahwan et al. 2003). Its basic idea is to allow more information to be exchanged between the parties. Thus the parties can, for example, offer a critique of their opponent's latest proposal, explaining why it is unacceptable or what they would like to see. This sort of information directs the negotiation to possibilities that are likely to be acceptable to both parties. The approach also allows the agents to explain why the other agent should accept their proposal. They can, in other words, try to persuade their opponent to change their region of acceptability. Moreover, it is possible to extend the negotiation to completely new areas, to change what is negotiated on. For example, if two parties have difficulties in finding acceptable terms of agreement for the purchase of a certain item, the seller might offer a free insurance or guarantee to the buyer.

[11] This is because the seller has decreased its price by 50% in five turns and therefore, the average decrease has been 10.

[12] The requirements for part II of the thesis are more complex and we will discuss them later. However, also in part II, our main focus will be on other questions and we will again be using (mostly) heuristic tactics (see section 2.3.2).

The downside of an argumentation-based approach is that it is quite complicated. An argumentation-based agent must be able to assess its opponents' arguments and form its own arguments to persuade the opponent. The more sophisticated the arguments that are allowed, the more complex these tasks become. In this work, we are interested in many issues in addition or instead of bilateral negotiation and using an argumentation-based approach would distract us from those other issues. It also does not bring anything that interesting to our problem because we will only have price negotiations. Thus, we do not build upon this line of work in the thesis.

2.2.2 Dealing with Changing Circumstances

We now turn our attention to commitments, decommitments and how they affect the common good. In most automated negotiation systems, the decommitment problem is not addressed at all (Sandhom and Lesser 2002).This means that, in effect, contracts are assumed to be absolutely binding: after one has been made, neither party can withdraw from it. However, in dynamic markets with incomplete information, this may not be optimal, since one or both parties may experience:

a. *changes in circumstances*: if the contract is not performed immediately after it has been entered into, the circumstances of one or both parties can change between the time of contract and performance:

   - the customer's requirements may change and affect the value of the service
   - the provider's circumstances may change and affect the costs of producing the service.

b. *changes in market situation*: if there is a delay between the contract and the performance, the market situation can change during this time:

   - the demand can increase (for example due to a large number of consumers entering) and this can increase the market price for the service significantly.
   - the supply can increase (for example due to a large number of providers entering) and this can decrease the market price for the service significantly.

c. *problems in interconnected negotiations*: the value of this contract may be fundamentally dependent on the outcome of other negotiations:

- the consumer may be expecting to use the service as a part of a larger service he is planning to sell to other entities, but he may fail to secure a contract or the price may not be as good as he expected.
- the provider may be planning to buy some parts of the service from other providers, but he may fail to secure these services or these services may prove to be more expensive than the provider expected.

d. *multiple contracts*: either party may be negotiating on the same service with multiple opponents and many of them can accept his latest offer, leaving him with multiple contracts even if he needs only one.

For example, consider a situation where agent A is negotiating with agent B on delivering a service s_1. In order to provide s_1, A needs a subservice s_2. A is unable to do this service itself, so it is therefore negotiating with C on its provision (case c above). The problem is now that if he manages to get a result in one negotiation, the other may still fail: B may buy the service elsewhere or C may sell his service to elsewhere. In the first case (B buys elsewhere), service s_2 has no value to him and in the second case (C sells elsewhere), he is unable to provide s_1. If we allow parties to choose from many contracts, the acceptance alone may not be enough to form a contract. So, A may accept both B's and C's offers and still be left without one or two of the contracts.

In a dynamic (even remotely realistic) environment, commitments cannot be absolute. Hardware failures, network problems and many other things may mean that despite the best efforts, the commitment cannot be fulfilled. On the other hand, commitments are useful, since they make it possible to make plans for the future. Thus, we need to find a balance between the two. We will now discuss shortly two possible ways to cope with the problem: contingency contracts (section 2.2.2.1) and renegotiation (section 2.2.2.2). We will then conclude by discussing our method of choice, the leveled-commitment contracts (section 2.2.2.3)

2.2.2.1 Contingency Contracts

The simplest solution would be to allow A to make a simple *contingency contract* (Raiffa 1987) with either B or C (as discussed in section 1.1.3.1). In such a contract, A would agree with B to provide s_1, *if and only if* he can get s_2 or A would agree with C to buy s_2 *in case* it manages to sell s_1. In theory, we could make our contingency contract explicitly state what we are supposed to do in every possible situation that might arise after the agreement. The main advantage of a contingency contract is therefore theoretical. If all possible contingencies can be covered and easily proved to have occurred, they remove the commitment (or replace it with another commitment or even another contract) every time a party fails to perform and, therefore, in case of trouble, the situation is always clear. In contrast, when a party to a binding contract fails to deliver, the original commitment still exists despite the failure to perform and the fact the commitment will never be satisfied. This opens potentially difficult questions, for example, on liability.

Another potential advantage of a contingency contract is that it makes the situation more transparent for all parties. For example, let us assume that A fails to secure the service s_2 and therefore is unable to provide the service s_1. No matter what the contract is like, B is left without the service s_1. The contingency contract does not change that. However, in case of a contingency contract, B knows the conditions under which A is not going to perform and can estimate the probability of them occuring or even notice them occuring (for example the only provider of service s_2 stops operations). This is not the case if B does not know what circumstances are essential to A's performance. Of course such transparency would mean that A would have to give quite a lot of information on its processes and sources to the outsiders.

However, the main problem of contingency contracts is that listing all possible contingencies can get very complicated and in any sufficiently complex system, there is always a possibility that something that nobody thought of (in advance) occurs anyway. Moreover, for contingency contracts to work, one must be able to prove the existence of each contingency (the failure to obtain service s_2 was not due to A's feeble attempts at securing it, but that there were circumstances beyond its control in play, for example), which might be difficult.

2.2.2.2 Renegotiation

As already mentioned in section 1.1.3.1, one way to remedy this problem of unforeseen events would be to *renegotiate* the contract, whenever something, which significantly changes the situation, occurs. Renegotiation takes place ex post, after the change, so parties know exactly what the new situation is and can possibly find a new balance between them. Renegotiation allows the parties to react to unforeseen contingencies when contracts are incomplete. Renegotiation may therefore be socially useful (Salanié 2000). Indeed, renegotiation is a recommended approach in both the Principles of European Contract Law and the UNIDROIT Principles of International Commercial Contracts in case of changed circumstances or hardship (case *a* above).[13] In both cases, if a change in external circumstances has significantly changed the balance between the parties, renegotiation is one way to restore it. Thus, if the other party is still interested in the performance, he might, for example, pay more or accept the delivery later.

However, the problem with renegotiation is that it relies on cooperation. A party can always reject any offer to renegotiate and is under no obligation to accept any changes to the contract. It is therefore difficult to ensure that an efficient outcome is always reached. To illustrate this point, consider the following example. Let the contract price be 10 and the cost of producing the service 8. Now, due to the changes in the environment, assume that the value of the service has decreased from 12 to 9. Thus the contract is no longer in the consumer's interest (since $9 < 10$), but it would still be beneficial to the society, since the benefit is larger than the cost ($9 > 8$). The parties could renegotiate the price to anything between 8 and 9 and both parties would be better off than by letting the contract go. However, the provider is not under any obligation to accept the lower price and can insist on the price of 10 instead. Now, in international trade the renegotiation often works because the parties will want to continue their business relations and it does not hurt to be fair every now and then, since the next time the roles might well be reversed.

There are, however, other situations in which a successful renegotiation is simply not possible. To illustrate this, let the value of the service decrease to 7 instead of 9. Now, the costs outweigh the benefits ($8 > 7$), so the contract is detrimental to welfare. Thus there is no price that both parties can agree on, so any renegotiation will fail. The best outcome would be to walk away from the contract. However,

[13]Principles of European Contract Law are created by a group of leading legal scholars from different European Union countries and they are intended to be applied as general rules of contract in the European Union (article 1:101). They cover all kinds of contracts. In contrast, the UNIDROIT Principles of International Commercial Contracts are created by United Nations' International Institute for the Unification of Private Law (UNIDROIT) to be used in the international trade between commercial entitites. They complement the United Nations Convention on Contracts for the International Sale of Goods (CISG) and other international conventions on international trade.

the provider will be able to make a profit with a price of 10 ($10 > 8$) and it has no obligation to withdraw from the contract, so it may well insist on it.

In theory, we could insist that the provider renegotiates in the first case and withdraw in the second case, but the problem is that in practice (with incomplete information) it may be difficult to say which of the two cases any particular situation is. In addition, it may be uncertain what the effect of a particular event actually is to the consumer's valuation of the service. Thus if the consumer knows that the provider will decrease its price to 8.5 if his valuation decreases to 9, he might claim that his valuation has decreased even when it really has not. On the other hand, it might be very difficult to prove the valuation really has decreased. Some of the events that cause changes in valuation may be very difficult to prove reliably with reasonable costs (for example the change in demand or in the buyer's preferences or goals). And obviously the consumer never has an incentive to indicate any changes that increase the value. Therefore the provider may not believe the consumer's arguments on a changed situation and may still insist on the higher price.

2.2.2.3 Leveled-Commitment Contracts

Given the analysis above, we believe that neither contingency contracts nor renegotiation are sufficient for the problems we are facing. They might offer a workable solution if we assumed that any change in circumstances would render the contract otherwise void, but the parties could not negotiate with anyone else on the subject for a certain duration (lock-out agreement). In this case, the parties might well have sufficient incentives to negotiate seriously with each other. Parties could also agree to release each other to negotiate with other parties.[14] However, this is still somewhat unsatisfactory. The third option to manage changing circumstances is to allow parties abandon their commitments, to decommit from them whenever they want. Now, of course if decommitting a commitment was always completely unrestricted, there would be no commitment. However, there is a way to make this work and it is the *leveled-commitment contracts* (section 1.1.3.1), which:

- allow unilateral decommitting for both parties at any time, but

[14]Another approach would be to deem the contract as one of risk distribution: each party would be responsible for any changes that make the contract less favourable to them. In this case, the contract would be binding unless both parties agree to let it go or in case the performance is physically impossible. This view is accepted by the English common law.

- require the decommitting party to pay a monetary fee (called a decommitting penalty) for doing so.

In this mode, the reason for decommitment is not relevant, any reason will do and the reason can be kept private (Sandhom and Lesser 2002). The name leveled commitment contract comes from the fact that the size of the decommitting penalty specifies the party's level of commitment. According to Sandholm and Lesser, the leveled commitment contract makes it easy to manage interconnected contracts and distributed search, they save time and computation[15] and allow reallocation of risk.[16] We think that the leveled-commitment contracts are the best approach to manage changing circumstances (requirement **R2**) and the necessary decommitments. However, this leaves us with a question about the optimal size of decommitment penalty and that is what we will discuss next.

2.2.3 The Decisions the Parties Make

The obvious question on the optimal size of decommitment penalty is not an easy one to answer. It depends on the circumstances and in ways that may not always be obvious at first glance. Generally speaking, one can say that if the penalty is very large, the situation is similar to cases where decommitment is not allowed at all. On the other hand, if the penalty was zero, there would be no commitment at all.[17] It is therefore clear that the optimal level for these penalties is usually somewhere in between, but the agent literature so far has concentrated on things that can be considered to be of secondary value. For example, Sandholm et al. (1999) have considered a problem of strategic decommitment, a situation where

[15]This may actually be quite problematic. According to Sandholm and Lesser, the ability to decommit saves computation, because the agents can bid on contracts using only approximations of contract value and they need not perform feasibility checks. In this case, only the offeror who actually wins the contract makes these calculations in detail and can decommit, if necessary. However, it is obvious that this can lead to very inefficient (initial) allocations of tasks. The winner is not necessarily the best provider, but is the one that underestimated its costs by the most.

[16]In market conditions, the contract price itself can be seen as a risk-allocation device. At the moment of agreement both parties have preferred the contract price to the uncertain market price of the future and have taken the risk of price changing to the direction that is bad for them. Allowing parties to decommit just because it later turns out they made a bad bargain changes this risk allocation. This can be seen as a good thing by some, but there is also a problem. If the market is very volatile and decommitment penalties small, it is very likely that one of the parties will decommit before the performance. This means that after the contract has been formed, the parties need to continue negotiating with other opponents, since the contract they already have is unlikely to be actually performed. This can be seen as wasteful use of resources.

[17]In case of a readily available market for the service, even a slightly better deal offered by another opponent would give a party an incentive to decommit.

one agent is not decommitting, although the contract is no longer useful to it, because it expects the other agent to decommit.[18] This may lead it to staying in a contract even when neither party really wants to.[19]

For us, the real issue is how the level of decommitment penalty affects the behaviour of the parties at a much higher and more general level. Now, contracts are a prevalent form of commitment in the society and they have been investigated thoroughly in many areas of science, including law, sociology, and economics. For us, the most interesting areas in this regard are law and game theory. In game theory, the problem of contracts has been investigated especially in the sub-field of law and economics, which investigates legal rules and economic efficiency. Originally one of the major areas of investigation was so called efficient breach theory, which investigated circumstances under which it was more efficient (from the society's point of view) to abandon the contract than to fulfil it. This was later named the *performance decision*. As we will see, it was shown that the legal rules in many jurisdictions about contract breaches can be seen to follow the approach this analysis prescribed.

However, this was not the end of the analysis. In the subsequent decades, the work has expanded and many other decisions were recognised. The level of damages (in legal terms and decommitment penalties for us) was shown to affect many decisions that parties make. Different authors classify these decisions slightly differently and identify different numbers of them. For example Craswell (2001) identifies no less than seven different decisions and says that even his list is not exhaustive but that any decision taken by the parties in relation to the transaction is, to some extent, affected by the extent of the liability they and their opponent bear. His examples include, among others, choosing who to enter contracts with and at what price, how much time to spend on looking for better alternatives before committing, how much effort to put into performing properly, how much to rely on the promised performance, how carefully to assess the risks involved and

[18] If the other agent decommits, the agent not only saves the decommit fee, but gets the opponent's decommitment fee.

[19] Sandholm et al. (1999) show how this can be avoided by calculating the optimal decommitting penalties in the following simple game. The game has two rounds. In the first one, the provider and consumer negotiate a contract and in the end either both accept it or it is rejected. In the second round, both parties get an outside offer from a known distribution (different for different parties). They then decide whether or not to decommit from the first contract (if it was made). The game itself is somewhat simplistic and assumes quite a lot. In particular, the only change comes from the counteroffers, thus both parties know the distributions the counter-offers are drawn from, that the values are statistically independent and that the distributions and other information are assumed to stay the same for the duration of the game. The game is obviously too simple for our purposes.

how much information to give to the other party. In other words the damages do affect the optimal behaviour of the parties in many ways, some of which may not be instantly obvious. And the problem is that these different decisions may prescribe very different levels of damages (decommitment fees) and in most situations many of these decisions happen in every transaction. This may mean that an economic analysis of a real-world situation may be very complicated.

In this section, we review the literature for the following four decisions in more detail: (Smith 2004; Craswell 2001; Kornhauser 1986):

- the performance decision (whether or not to perform) (section 2.2.3.1),
- the reliance decision (how much to rely on the opponent's performance)(section 2.2.3.2),
- the contract decision (whether or not to enter the contract at all)(section 2.2.3.3), and
- the selection decision (who to negotiate with)(section 2.2.3.4).

These four decisions were chosen because they are the most important ones in this setting. The performance decision is much discussed in the literature and the legal rules on contracts can be, in many cases, justified with analysis of this decision. The reliance decision is the flipside of the performance decision: when one party makes a decision whether or not to perform, the other one makes a decision whether or not to rely on the performance. The contract decision is obviously central to us, since we are interested in the changes of circumstances (requirement **R2**) and one way to avoid decommitments is to consider whether or not to enter a contract in the first place. The selection decision is also an obvious choice, given that we have heterogeneous participant populations (requirement **R3**). Because many of these decisions have not been discussed in the context of agent marketplaces and because it is unclear how big the differences between different policies are in different situations, we will investigate each of the decisions separately.

2.2.3.1 The Performance Decision

In the *performance decision*, there is an existing contract and the parties to that contract decide whether or not to perform according to it. The contract party makes a decision between staying in the contract and decommitting from the contract and paying the decommitment fee. He will choose to decommit only if the decommitment actually improves his utility more than the fee to be paid reduces it. In most of the literature, the decommitment fee is therefore mostly seen as a deterrent of decommitment. This view, however, overlooks the effect a decommitment has on the other party (the victim) and on the society in general. In particular, the decommitment will usually decrease the utility of the victim, because he will lose the profit he was expecting. Moreover, it is possible that he has already accrued some costs (preparing for its performance) before decommitment occurs. When the contract is abandoned, these efforts may become useless. Now, if these lost profits and accrued costs outweigh the benefit the decommitter receives from decommitment, the decommitment actually decreases the sum of utilities of the parties and is, therefore, detrimental to the welfare of the society as a whole.

In general, the computer science literature has not discussed these issues at all. However, Andersson and Sandholm (2001) have recognised the possibility of using the price of the contract as a basis of the decommitment penalty, but instead of full compensation they suggest that the penalty should be selected as a percentage (or a more complex function) of the price. Another way they suggest is to make the penalties compensate the victim of the breach for its lost profit. If combined with the compensation of useless work, this is essentially the expectation damages.[20] In this vein, Nguyen and Jennings (2004) have suggested a dynamic decommitment fee, which is calculated as a percentage of the (buyer's) utility of the deal.[21] Attaching the decommitment penalty of both parties this way only to the buyer's utility may be problematic, since obviously the fee is not connected in any way to the seller's utility[22] and it does require the buyer to reveal his utility

[20] Andersson and Sandholm point out that because the victim has an incentive to lie about its profit, some sort of mechanism would be required to calculate the loss. The state of both agents might have changed since the contract was made, so the expected lost profit may have changed between contracting and breaching time. In the extreme, the lost profit for the victim can be negative at breaching time, that is, also the victim of the breach benefits from being freed from the contract obligations.

[21] They experiment with different starting and ending levels and seem to get a result that the higher decommitment fees mean lower utilities. However, it seems that their services cost only at the time of performance, so the problems discussed here are not presented in their case. In addition, the starting and ending levels of their tests are quite close to each other (0-0, 5-10, 10-10, 10-20, 25-25, 25-50, 50-50, 50-100, 75-75 and 100-100) and the formula they use to calculate the costs is linear in time. The sellers in their model also do not necessarily make their decisions on decommitment rationally, but they may be loyal (never decommit) or partially so (decommit with probability of 0.5, if a better contract is found).

[22] The decommitment fee therefore depends on how well the buyer has succeeded in negotiations. If he did badly and his utility is low, the decommitment fee is also low. On the other hand, the utility of the seller is probably relatively high, so the decommitment fee is not compensatory at all if the buyer decommits.

to the seller.[23]

Also Excelente-Toledo et al. (2001) have suggested variable penalty contracts as an extension of levelled commitment contracts, that make a connection between the actual costs and the decommitment costs. They offer two penalty schemes that are especially interesting to us, partially sanctioned and sunk costs. In the former, the actual fee depends on the state of the coordination activity, its participants and the estimate of the profit. However, the amount paid is agreed in the contract and the compensation may therefore be only partial. The latter compensates the effort that parties have invested so far but not the expected profit.

All schemes discussed in the last two paragraphs aim to compensate the damage the decommitment causes to the opponent, at least to some degree, which we think is the right approach. However, none of them is entirely consistent (i.e. they do not consider different relevant factors, such as actual costs and profits) and we think that they do not go far enough (full expectation damages). We will now discuss the issue of decommitment in terms of contract law, to see if it could provide us some ideas how to proceed with decommitment fees.

In the law of contract, the contracts are binding (pacta sunt servanda) and any kind of non-performance is a *breach* of contract. When there is a breach, the law of contracts gives certain remedies for the opponent. There are two main ones:

- A *specific performance* means that the obligations laid down in the contract are enforced. If a person promised to give something, the court will order that something to be given. A court order can be enforced by other officials and, in extreme cases, even force could be used. The aim is to make the parties perform as they had agreed to do in the contract.
- The *damages*, on the other hand, are a monetary sum that aims to compensate for the lack of performance. The expectation damages are usually awarded.

The different legal systems differ as to which remedy is dominant. In some countries, the specific performance is the main remedy and the damages are usually used only if the specific performance is not possible. Other countries allow the other party to choose which remedy to invoke. In common law countries (roughly

[23]Such information may weaken its negotiation position in the next negotiation between the parties.

the same countries that have English as an official language[24]), however, the damages are preferred and the specific performance is used only in specific situations, where damages are deemed to be very difficult to assess.

Now, in a common law system, a party to a contract can get out of most contracts simply by paying these damages. This idea is not unlike the decommitment penalty discussed before, but there is one big difference: the damages are *compensatory*. That is, they try to compensate the opponent for the damage he suffers due to decommitment, where the decommitment penalty usually is not.[25] This means that if non-compensatory decommitment penalties are used, the selfish parties worry only about their own welfare and can decommit in situations where the benefits they get from decommitting are smaller than the losses suffered by the opponent (Sandhom and Lesser 2002). This is usually detrimental to the welfare of the society. The solution offered by the law and economics literature is easy enough: set the decommitment penalty (damages) equal to the value of performance. In this way, the opponent's losses are internalised to the parties' decommitment decision and the decommitment occurs only if a party expects that the benefits are larger than the losses (i.e. only when decommitment increases welfare).[26] This is exactly what the contract law rule on damages can be seen to be doing.

Generally the goal of contract law is to facilitate contracting in situations where the agreement and performance or the performance of different parties do not occur at the same time. This is exactly what we have in our setting: We have costly and time-consuming service preparation (requirement **R1**). The seller must begin preparation of its performance (the service) before the buyer make its performance (the payment). The seller may have to, for example, buy the necessary materials, compile the data and start the analysis. If these actions cost money and are useless if the service contract is decommitted before the performance, it is clear that a

[24]However, this classification is not precise and should not be taken as such. For example, Scotland's legal system is not based on common law.

[25]According to Sandhom and Lesser (2001, p. 216): '...explicitly allowing decommitting from the contract for a predetermined price is used as an active method for utilizing the potential provided by uncertain future.' This means that the decommitment penalties are often very low and too low to be compensatory. As we explained earlier, there are some papers that suggest a more compensatory approach (e.g. Anderson & Sandholm and Excelente-Toledo et al. discussed in section 2.2.3), but, in our view, they do not go far enough.

[26]However, Sandholm and Lesser think that the results of law and economics cannot be used, because they discuss the level of damages *after* there is a breach. However, this is a misunderstanding of sorts. Any one-sided decommitment before the time of performance is dealt with under the rules concerning *anticipatory breach of contract*. Thus the contract is binding and cannot be unilaterally terminated. If a party to a contract informs the other that he is not going to perform the contract, this decommitment is not usually effective unless the other party accepts it, but even in this case the victim is entitled to damages.

reasonable decommitment penalty should compensate for these losses. Moreover, the prospect of such compensation in case of trouble allows people to take these actions without needing to worry about the costs (they will be compensated in any case as long as the buyer is solvent) and other legal rules allow people to pay in advance, since in case of non-performance, they can usually get their money back. And this allows parties to trade even in the face of uncertainties about the future and the other party.

Also one could also consider that legal rules have also evolved during hundreds of years of trading, and therefore one could argue that they should represent a reasonable starting point for the rules of any trading system. Any situation that can happen in a trading context has happened countless times before and legal rules are likely to cover it. Of course the different legal systems do vary in details, but the big picture is usually surprisingly similar in most legal systems in the world.

Since our setting is just what the legal rules are about and the rules themselves seem to facilitate efficient trading between the parties, it seems a very promising approach to see how the law manages more complicated situations and how the legal rules can be used to adjust decommitment policies in these more complex settings. That is exactly what we will do later in chapter 4. Since most of the literature and other work on these decisions, is on the performance decision, that will be our main focus too. We will discuss many different settings and develop decommitment policies for them and investigate also settings with incomplete information.

The law and leveled-commitment contracts are not a perfect fit, however. As discussed above, with the leveled-commitment contracts, the reason the party wants to decommit is not relevant: the party can decommit at any time for any reason, but on the other hand, it needs always to pay the fee. The law, however, has traditionally in most countries seen differences in the situations: Some reasons are considered to be better than others and there is usually good explanation for such exceptions. The biggest difference here is that in law, there are special situations where it might be possible that the party's duties end without any payments or can end for less than full damages. These, obviously exceptional, situations are handled in English law under the doctrine of *frustration*. It deals with situations in which the circumstances change between the time of agreement and performance, in such a manner that the performance, as agreed, is impossible (for example a specific item sold is destroyed), illegal (for example handguns are outlawed before the seller can provide the buyer with one) or has changed to something completely different (a trip to an area that has become a war zone). In addition, frustration requires that the problem could not have been foreseen or won with reasonable costs. If frustration is permanent, the contract is void. This means that the obligations of the parties no longer exist and any performances (or at least the equivalent value) already made should be returned.[27] In an agent-based market an impossibility to perform could mean for example power outage, network or other hardware malfunction or something similar.

However, since these situations are (by definition) exceptional, the way they are handled may not have a huge impact on the whole. The fairness aspect (if there was nothing a party could have done better, having to pay a fee may seem unfair) might justify using these exceptions also in agent systems. However, there are potential problems. On one hand, if these types of exceptions are allowed, it could create incentives to make fraudulent claims and to counter that one would need a process for evaluating the claims that the exceptional circumstances are at hand and that may be tricky in some situations, especially the 'could not have been anticipated or avoided' part. This would make the system much more complicated than 'you fail, you pay'. And on the other hand, in different legal systems, these exceptions are defined in different ways and they are somewhat different in extent and applicability. In an international marketplace, one would need to clearly explain what the rules in this particular market are and exactly when would exceptions be granted. In this regard 'you fail, you pay' or 'no fees' would be much simpler. Due to these reasons, we will not consider these exceptions in this work.

2.2.3.2 The Reliance Decision

If the performace decision was the only relevant decision, things would be easy. To maximise common good, one would set the decommitment fee equal to the opponent's losses and it would only be a problem of getting that fee close enough

[27]The applicability of this doctrine is very limited under English law. In many other jurisdictions there is a similar but usually a wider doctrine of *force majeure*, which allows non-performance without damages in case of war, strike or other similar circumstances that could not be foreseen, won or bypassed with reasonable costs. Once the problem is removed, the performance must be made. However, the opponent is usually in these situations given a right to withdraw from the contract without any damages. Another doctrine of (economic) *hardship* may allow a free decommitment or one with partial damages, if the cost of performance has increased significantly. This usually requires a very significant change that could not be anticipated or avoided. As explained earlier (in the beginning of section 2.2.2.2), renegotiation is usually a preferred means to deal with these situations.

to the actual losses under incomplete information. However, that is not the whole story. As mentioned, there are other decisions to consider.

The optimality of the full expectation damages decommitment policy assumes that there is already a contract (the price is set), that the losses of the other party are fixed and that this party is unable to influence these costs. However, in many situations this may not be the case, but on the contrary, the party may have a strong impact on its costs and expected profits. To see this, let us consider the situation from the other side: When a party makes a performance decision on whether or not it will perform itself, the other party makes a reliance decision to decide how much to rely on the other party's performance.

In more detail, the party shows its reliance on the opponent's performance by taking actions that enhance the value of this future performance. Such preparation might involve, for example, getting complementary services, such as booking a coach ticket to the airport in anticipation of an airline providing a trip somewhere nice or getting a good bottle of wine to accompany a nice meal that somebody has offered. The preparation could also mean making changes to or configuring one's equipment to best take advantage of the performance, like getting a manufacturing plant ready to make product A as soon as the raw material B arrrives. Also in a more abstract form, preparation might mean turning down alternative plans like not taking a job A because one is expecting to work on a job B at that time, and so on. Also the seller makes a reliance decision and can have plans for the renumeration he is expecting from the buyer.

Now, in many cases such preparations can be useful (beneficial reliance (Goetz and Scott 1980)).[28] For example, it is good to have a coach ticket in advance, so that one can be relatively certain to make it on time to the airport (no need to worry, if there is room on the coach). It can also mean that one can avoid more expensive last-minute options, like taking a taxi to the airport. A meal can be enhanced with a good wine increasing the value of the meal itself. Making the changes to one's own equipment in advance might save time and resources, because one can then set to work the second the required resources arrive, and so on. Reliance on the future performance is therefore often useful, to the relying party but also to the society in general. For example, it's good for the society that the expensive resources are in effective use, or that the services are combined with other services

[28]The detailed reliance model in (Goetz and Scott 1980) considers a non-reciprocal gift and is therefore not directly applicable to our problem. They do consider many relevant issues and other decisions, also in a reciprocal context and that is what we base much of the discussion here.

to enhance their value and sometimes the best way to make these improvements is to anticipate a future performance.

Of course the flipside of this coin is that if the other party fails to perform, the relying party might be worse off than if he had not relied on the performance at all (detrimental reliance (Goetz and Scott 1980)). Thus, he might have a coach ticket or a bottle of wine that he has very little use for (and may not be able to refund). Configuration or other changes might turn out to be useless and configuration or changes have to be redone (or even just revert to the earlier settings). He might not be able to get any work to the suddenly open slot in his schedule and so on.

Now, it's easy to see that the decommitment policy has a profound effect on the reliance decision. The higher the decommitment fee in case the opponent does not perform, the more the performance can be relied on. However, there is a level of reliance that is better for the society than some other levels (the optimal level of reliance) and it is not necessarily always the highest possible reliance. It can be counter-productive for the society, for an invidividual to to make extensive and expensive preparations for a performance that is unlikely to happen, because such non-performance only means costs to somebody from the society's point of view. The party that must shoulder the cost in the end depends on the decommitment policy.

In the performance decision, the trick was to internalise the opponent's losses to the decision-maker's (decommitter's) decision by making him liable for the opponent's losses. This made him decommit only when it benefits the society. The approach here is the same, but the prescription is decidedly different. The only way to make the party making the reliance decision to choose the socially optimal level of reliance is to make his compensation in case of non-performance independent of the reliance level chosen. This is because compensating for additional reliance would encourage the party to rely over-optimally on the other party's performance.

To see why this is so consider the following situation (from Goetz and Scott (1980)). A party to a contract knows that the opponent's performance is less than certain (probability of performance w_s less than 1) and let $U_{performance}(\rho)$ denote the utility in case the other party performs according to the contract (beneficial reliance) and $U_{decommitment}(\rho)$ denote the utility if the other party decommits instead (detrimental reliance) without any fees. The ρ is the level of reliance. The $U_{performance}$

is increasing and $U_{decommitment}$ is decreasing in ρ to the maximum level of ρ_{max}.[29] This means that by increasing reliance, a party can increase $U_{performance}$ but pays for that with a decrease of $U_{decommitment}$. Now, the expected utility in the situation is

$$EU = w_s U_{performance}(\rho) + (1 - w_s)(U_{decommitment}(\rho) + f).$$

An interesting point to make here is that the optimal reliance level (ρ^*) is independent of the fee f as long as f does not depend on ρ. This follows because the first derivative of EU in relation to ρ does not contain f at all. In contrast, it is equally clear that compensating for the (extra) reliance would be a very bad policy. For example, the full *Expectation Damages* policy where all reliance would be compensated for would always lead to the maximum possible reliance ($\rho = \rho_{max}$).[30] So, to make an extreme example in a contract of certain type of bolt, it would make sense for the buyer to keep a hugely expensive machine, like a paper mill, running although it might be completely destroyed if an order of one bolt does not arrive in time, because the buyer of the bolt is indifferent between the bolt arriving and the machine continuing to make paper and the bolt failing to appear and the machine getting destroyed, because the bolt seller would cover the difference. Of course such a result would be non-sensical and no court would ever order such compensation to be made. The law has its own mechanisms for limiting compensation.

However, a problem remains. Usually in the literature the optimal policy in case of non-performance is said to be 'zero' compensation (Smith 2004; Kornhauser 1986) and the model above seems to indicate that any reliance-independent fee (also non-zero) will do equally well. There is no conflict here: zero simply refers to the reliance that should be covered. This can perhaps be better illustrated by considering the seller. His costs come from reliance: by undertaking costly and time-consuming preparation he relies on the performance that the buyer has promised, because only through that will he be able to cover his costs and get a profit. If he does not start preparations, he will not incur any costs but then he will not be able to reap any benefits later either ($U_{performance} - U_{decommitment}$). Now, the analysis on the reliance decision prescribes that none of this reliance

[29] For values $\rho > \rho_{max}$, we could say that both $U_{performance}$ and $U_{decommitment}$ are decreasing. The level ρ_{max} therefore denotes the highest possible reliance that makes sense, i.e. the increase in value is greater than the cost. After this value, the costs are greater than the increase in value and even if performance was certain, it was not rational to rely that much.

[30] This is because the fee f would be equal to $U_{performance} - U_{decommitment}$ and therefore the party would get $U_{performance}$ with probability of 1. From the above discussion it follows that $U_{performance}$ is maximised when $\rho = \rho_{max}$.

should be covered by the decommitment fee, because if it is covered, the seller has no incentives to consider trust-worthiness of his opponent and if his costs and profits are always covered, he can start working even if he knows the buyer is very likely to decommit at some point.[31] Only a performance will benefit the society. Decommitments will only bring costs and possibly re-distribution of wealth. This of course contrasts with what we learnt in the performance decision, where the parties covering for each other's costs and profits were essential. We will (in chapter 5) build on this model and investigate how different decommitment policies (rules on decommitment fees) will affect the common good.

2.2.3.3 The Contract Decision

The party to a potential contract makes a *contract decision* when it decides whether or not to enter into a proprosed contract now. An offer that seems lucrative now might lead to a counterproductive, even disastrous, contract if the party's circumstances or the market situation adversely change before the contract is to be performed. The higher the required level of commitment (the higher the decommitment fee) and the higher the probability of such adverse changes occuring, the less inclined a rational agent should be to enter into a contract in the first place. From the society's point of view, very risky, high cost contracts may well be inadvisable.[32]

To this end, Smith (2004) has argued that in relation to the contract decision the expectation damages are sometimes too high. Instead, reliance damages that include the foregone profits of other possible contracts (the profit that could have been achieved outside this particular contract) would cover the other party sufficiently (ensuring non-negative profit) and sometimes facilitate contracts that

[31] In case the *Expectation Damages* policy is used, he is even indifferent between the case in which the buyer wants the service and the case where he does not.

[32] Of course in some cases such contracts or projects might be in the society's interest, if the payoffs for success are sufficently large. For example, trying to build the world's first fusion reactor is going to take a long time and cost a huge amount of money and other resources even in the most optimistic scenarios. And still, the success is far from certain. But, on the other hand, a clean, almost unlimited source of energy would obviously be *very* useful for any society. Given the costs and risks involved, some major governments (including European Union, USA and Russia) have decided to cooperate in a joint project in this area, see `http://www.iter.org/`.

would not be possible if full expectation damages were offered (Fuller and Purdue 1936).[33]However, in an efficient market, the difference between the two is very small or even non-existent because in such a setting, there would be many providers selling a similar product for a similar price. Therefore by taking on one provider, the buyer has in effect lost an opportunity to make a very similar profit with some other provider. Therefore the loss is roughly equal to cost + profit (i.e. the expectation damages). However, in a situation with fewer outside opportunities, the second-best option can be much worse, in the extreme case, it might be zero (if there are no viable alternatives). This might happen, for example, if the provider has a lot of spare capacity that he would not have any use for outside this one particular contract.[34] Compensating the expectation damages automatically does keep the seller's risk to a minimum (any contract that will have a positive utility in case of success will have that also in case of the buyer failure) but it will also mean that sometimes the buyers will not enter into a contract even if it was in the seller's and in the society's interest. In other words, a potential buyer might in some cases be dissuaded from using the service if, in case of decommitment, he would have to pay for the seller's profits that the seller would never have had a chance to get without the buyer in question.

A numerical example might better illustrate Smith's idea. So, assume that the cost of providing a service, which has to be paid immediately, is 5, the price of the service is 7 (due to the low season) and the value of the service to the buyer is 13, and the probability of the buyer needing to decommit is 50%. Now, the buyer's expected utility in the cases with reliance and expectation damages would be:

$$\begin{aligned} EU_b(\text{contract} \mid \text{reliance damages}) = &\ 0.5(13-7)+0.5(-5)=0.5 \\ EU_b(\text{contract} \mid \text{expectation damages}) = &\ 0.5(13-7)+0.5(-7)=-0.5 \end{aligned}$$

and for the seller:

[33]Smith's proposition means compensating so called opportunity cost (or best alternative to negotiation agreement, BATNA) instead of expectation damages where the profit compensated is the one the victim expected to make from the existing contract. The legal rules are mostly interested in the performance decision and ignore these other decisions. Legally, entering into a contract means accepting a liability for the opponent's reasonable profits from the contract also in case of non-performance. This holds also when substitute contracts are used. This has nothing to do with the opportunity cost.

[34]This might be for example because his business is very seasonal and outside the season there is very little demand.

$$\begin{aligned} EU_s(\text{contract} \mid \text{reliance damages}) = & \quad 0.5(7-5) + 0.5 * 0 = 1.0 \\ EU_s(\text{contract} \mid \text{expectation damages}) = & \quad 0.5(7-5) + 0.5 * 2 = 2.0 \end{aligned}$$

So the seller would want this contract to happen. If we assume that the value of the buyer's outside option is zero, he will choose to enter the contract when the reliance damages policy is used, but not enter into a contract when the expectation damages policy is used.[35] For the society, the benefit of the contract would be in both cases $0.5(v - c) + 0.5(-c) = 0.5(13 - 5) + 0.5(-5) = 1.5$ and therefore the contract would be also in the society's interest.

However, this is not the whole story. It misses the important point that the contract price is a major risk allocation tool and if the decommitment policy is clear and all the relevant risks are known by both parties in advance, the parties can use the price to find a mutually acceptable balance of profits and risks with almost any policy (Goetz and Scott (1980)). To see how this works in practice consider the buyer in the numerical case above. In cases where the expectation damages policy is in use, he is unwilling to enter into a contract because the risk of decommitment and the fee are too high. Now, for the seller the high decommitment cost is a good thing: It increases his expected utility from 1.0 to 2.0 and, therefore, he may well be willing to offer a lower price. This lower price, in turn, increases the utility of the potential contract to the buyer and makes it more lucrative. So, in the numerical example above, under the expectation damages policy, the seller might offer a price 6 for a buyer that is unwilling to go for it with the price of 7. This would make the expected utility for the buyer equal to $0.5 \cdot (13 - 6) + 0.5 \cdot (-6) = 0.5$, which would be enough to persuade the buyer to enter into the contract. The contract price can be used very efficiently in this way.

However, sometimes this mechanism is unable to find a good balance and sub-optimal contracts might be entered into or parties might not be able to agree on a price although a mutually acceptable price exists. For the price mechanism to work properly, the parties must be able to assess relatively accurately the probabilities that they or their opponent will perform. If no such information is available,

[35] Of course the buyer would be even more willing to get into these contracts if, for example, the fee would be zero and also the society might be better served in some lower level of damages, because there might be cases where a contract could be reached with a lower level of compensation when it is not possible with reliance damages. For example, if the probability of failure is 60% instead of 50, the buyer would not enter under reliance damages (expected utility -0.6) but a compensation level of four would make the buyer indifferent between entering into the contract or not and the society would still benefit, its benefit from this would be 0.2. However, here, Smith seems to think that covering the 'victim's' costs (ensuring him non-negative utility) is fair under the circumstances.

it may be impossible to find an optimal contract price. Also the probability of performance in many situations depends on the fee (the higher the fee the more probable the performance), which may make this assessment more difficult if the fee varies. We will investigate this price adjustion mechanism and what it means when we have none, one or both parties using it in different settings.

We will concentrate our efforts in this work (in chapter 6) in investigating this price adaptation to different decommitment policies. We will use different decommitment policies and allow none, one or both parties to take the risks of later decommitments into account and see how that affects the common good.

2.2.3.4 The Selection Decision

Before a party even starts a single negotiation, it needs to make a selection decision. That is, decide who to negotiate and potentially transact with. The selection decision is often mentioned as one of the relevant decisions, but it is rarely discussed in detail or if it is the models are often quite simple. This may be because it is often quite simple. For example, the selection decision might have to do with the opponent's reliance or quality of service. From the society's point of view, the situation is simple: A non-performance does not benefit the society but only causes costs that are (in part) away from everybody (for example, extra hassle caused by non-performance does not benefit anybody, also effort or materials may be wasted and so on). The performance, however, is usually beneficial to the society. This is why performance should be encouraged. This does not necessarily mean that the most reliable opponent should always be selected, but sometimes a small risk of non-performance may be acceptable if the benefits are high enough. In other words, from the society's point of view, the parties should select their contracting partners so that this:

$$EU = wU_{b+s}(\text{performance}) + (1-w)U_{b+s}(\text{decommitment}),$$

is maximised. This may seem very similar to the reliance decision, but here we do not have a contract. Rather we are only deciding who we should negotiate with in the first place. In contrast, in the reliance decision, we had a contract and decided how much to rely on it.

The interesting observation here is that awarding the victim full expectation damages makes him indifferent between performance and non-performance (the same utility in both cases[36]) and this may not allow a party to make a decision that maximises the common good. And what is even worse, if we have an over-compensatory fee (a fee that awards the victim with more than expectation damages) it means that the victim will actually prefer non-performance and choose its contracting partners accordingly. Interestingly, here the common-good-maximising fee would be zero or even slightly negative (decommitment bonus).

We will investigate (in chapter 7), how different decommitment policies affect the choices of the buyer that is considering only its own good and the impact these choices will have on the common good.

2.3 Concurrent Bilateral Negotiation Strategies

In the part II of the thesis, we will be interested in concurrent bilateral negotiation, and especially tactics and strategies needed to succeed in such environments. As we discussed in the introduction, in this part, we are interested only in the welfare of one of the parties, the buyer, and we do not generally consider common or the seller's good.

Here, in this literature review, we first discuss our reasons for choosing concurrent bilateral negotiation for our interaction model (section 2.3.1). We then proceed to discuss three levels that any concurrent bilateral negotiation model meant to negotiate on interconnected services needs to address. First, we discuss bilateral negotiation (section 2.3.2). We discussed the basic bilateral negotiation models already in section 2.2.1 and here, we will focus on how a bilateral negotiation model can be used to better manage negotiations in dynamic and open markets that have heterogenous players in them. Our second stop is the concurrent bilateral negotiation itself (section 2.3.3). We focus on the key issues of managing more than one negotiation on the same service. The third topic is managing interconnected negotiations, negotiations on different services that are somehow connected to each other (section 2.3.4). In such cases, the progress and results of a negotiation on one service can have impact on the negotiations on some other service.

[36]This assumes of course that the compensation in case of non-performance is perfectly reliable, so always happens.

2.3.1 The Case for Concurrent Bilateral Negotiation

In modern day commerce, negotiations are used rarely, especially when the service or product to be traded is of small value. The dominant way of making business, especially in the business-to-consumer relations, is that each service provider has its set of standard terms and it offers its services for one fixed price (*fixed posted pricing*). In these situations, the consumer's choice is limited to 'taking' or 'leaving' the offer as it stands (Atiyah 1995). Any attempt to modify these terms in a one-off transaction is usually turned down.

However, these fixed term contracts are not the perfect tool for either party. If the customer is not willing or able to accept the offered standard terms, he cannot use the service at all. For the provider, these customers are lost sales. He could obviously offer more than one set of prices and terms, but there might still be customers that cannot accept any of the options, but with whom the provider might be able to reach a mutually acceptable agreement in negotiation. The standard form agreement is essentially a way of saving the time, trouble and expense of making a contract (Atiyah 1995). The terms have to prepared only once and then they can be used in many thousands or even millions of transactions over time.[37]

Now, as already mentioned in section 1.1.2, information technology can decrease both of these costs dramatically. First, in online stores the price information is usually stored in a database and changing the price of any item or a group of items is very easy. In addition, the store owners can let a group of software agents (called pricebots by Kephart et al. (2000)) control the prices on their behalf. Another change caused by the rise of information technology is the possibility of collecting a significant amount of information from the consumers (Ancarani 2002). Together these two changes have made it possible for the providers to use *dynamic posted prices*, a modern-day variation of fixed posted prices. Basically this refers to take-it-or-leave-it pricing, in which the seller can change the price at any time (Kephart et al. 2000). The goal of dynamic pricing is to optimise the seller's profits by adjusting its prices to correspond to a customer's willingness to pay (Weiss and Mehrotra 2001). In its extreme form, the providers could use

[37] Although writing a good standard form agreement may (arguably) be more difficult (expensive) than writing a good individual contract, the cost per transaction is much smaller, when the number of transactions is high. Therefore the fixed price standard term contracts are usually the cheapest option in these circumstances (Rothkopf and Harstad 1994). Using the same terms with all customers also makes it easier (cheaper) to manage hundreds of contracts. The businesses do not have to spend time remembering (or the cost of recording and finding) what was agreed with whom.

the pricebots to customise the price for each market situation and even for each customer given the customer's earlier purchases, his estimated future value and price consciousness, prices of the competitors and other information. On the other hand, the consumers can also use information technology to automatically compare the providers, so the full implications of these changes are still uncertain.

Second, automated negotiating agents have the potential to radically reduce the costs related to negotiation. The agent only needs a computer to run on, a network connection to communicate with other agents and some electricity for the computer. One modern computer can easily run hundreds of agents, a reasonable Internet connection is relatively cheap (especially per negotiation) and the computer uses only very limited amounts of electricity. Automated negotiation is also often very fast. Even the most complicated automated negotiation is usually over in a matter of seconds. Thus transaction costs per negotiation are typically very low. These advantages make automated negotiation a viable option in many situations, where any human negotiation would be too expensive.[38]

Our case for using concurrent bilateral negotiation is four-fold. First, we will explain why the dynamic posted pricing we just discussed is not enough in many settings and why some form of negotiation is needed (section 2.3.1.1). Second, we will discuss why it is useful to negotiate with more than one opponent (section 2.3.1.2). Third, we will explain why we think these negotiations should occur (at least partially) concurrently (section 2.3.1.3). Finally, we will explain why we do not use auctions, although many people think that auctions are going to be a very important method in these electronic service markets (section 2.3.1.4).

2.3.1.1 Dynamic Posted Pricing versus Negotiation

The main problem of dynamic posted pricing is the fact that in the end of the day it is the provider who must set the price and other terms of his posted offer. Even with all the information that can be gathered on the customers, the picture will be incomplete and there will be situations and consumers that are difficult to manage even for the most sophisticated algorithms.

Now, negotiations can offer *two-sided interaction* between the parties. In such interactions, both parties can indicate what they want and what they are willing to give in return. This will alleviate the problem of incomplete information and

[38] Also any bookkeeping activities of parties could be easily (although obviously not without cost) automated.

can lead to better outcomes for both parties. The provider can get a better understanding of his customers' needs and use this information to make his selection more interesting to his potential customers.[39] In one-issue (price) negotiations, the negotiations can be used to find out what a given customer is willing to pay and at what price, a given provider is willing to sell. The importance of two-sided interaction would obviously be enhanced in multi-attribute situations, where a service has many attributes that can be configured interactively.[40] Although we will not have multi-attribute negotiations in this work, the general plan, of course, is to investigate concurrent bilateral negotiation in this simple setting so that our results might later be extended to the multi-attribute settings, where the concurrent bilateral negotiation can be even more useful.

2.3.1.2 Negotiating with One versus Many Opponents

Once we have decided to negotiate, we are then faced with the decision about how many opponents to negotiate with. Clearly, if the market has thousands of providers and we just select one of them more or less at random and negotiate with it, we run the risk of making a bad choice: the opponent we selected might not offer a quality level we would like, he may have a very high reservation price, or he may employ a stubborn negotiation strategy with us. If, on the other hand, we negotiate with multiple opponents, we get a much better view of the market and increase the probability that we manage to find at least one provider who is willing to offer us reasonable terms.

Negotiations with more than one opponent might also give the agent ideas on the market situation, (i.e. whether the current opponent is just being unreasonable or it is a seller's market). Since the providers may be able to collect a lot of information about the consumer population over time and can use that information to their advantage (Grover and Ramanlal 1999), the consumer must try to correct the imbalance by collecting more information themselves and, in the process, try to make the different providers compete for their custom. This should encourage the providers to be more flexible with their demands.

[39]Through negotiations he may also get information on why his offerings are *not* competitive. This information may be quite difficult or expensive to obtain by other means.

[40]Such configuration clearly requires two-way interaction between the parties, especially if there are too many possible configurations to list all the possibilities as separate services.

In short, it is clear that we should be negotiating with more than one opponent for each service. On the other hand, due to the limited negotiation resources (requirement **R4**), we may not be able to negotiate with all of the potential opponents in a large market. We therefore need to choose a subset of providers to negotiate with.

2.3.1.3 Negotiating Sequentially versus Concurrently

When it comes to negotiating with multiple parties on the same item or service, there are two main options: we can either negotiate with one of them at a time (sequentially) or many or all of them at the same time (concurrently). Sequential negotiations are easier to analyse mathematically and they have been used in some market models in game theory, usually with a random matching of buyers and sellers (for example Rubinstein and Wolinsky (1985) and Binmore and Herrero (1988)). Concurrent negotiation, on the other hand, is more difficult to analyse and there are no general game theoretic models that use concurrent negotiations as a market mechanism.[41]

However, in dynamic environments where anything can change at any time, concurrent negotiation has the clear advantage: it can get a good view of the market situation more rapidly. Sequential negotiating agents learn the possible offers only one by one and they have to turn down many offers and end many negotiations to get that view. In dynamic markets, by the time an agent is in a position to make an informed decision, the best providers might no longer be available or they may not accept the earlier price. In contrast, in concurrent negotiations, agents quickly learn the different types of offers in the market, can try different offer types with different opponents, and can relatively quickly recognise the most promising candidates and concentrate the negotiation effort on those. Concurrency allows a negotiator to use information from one negotiation in another negotiation while both negotiations are still running. This wider view can be particularly useful if, for example, the agent's requirements change completely or there is limited time. Thus concurrent negotiations are much more flexible and since flexibility is very important in dynamic environments, we prefer concurrent negotiations to sequential ones.

[41] The one exception to this is Chatterjee and Dutta (1998) who use only two sellers and two buyers. However, they assume that parties have complete information on each other's parameters, which obviously is not the case in the scenarios we consider (we explicitly assumed limited information **R6**).

As is so often the case, this flexibility comes with a price. When there are multiple negotiations going on at the same time, it is entirely possible that we end up with an agreement with more than one opponent at the same time. This may be a problem if we only need or have one item. Thus we must either be able to explicitly accept any agreement (our offers are not binding, but a separate acceptance by us is needed) or to decommit from the contracts we do not need.[42] Another problem is that concurrent negotiation requires more computing power and more communication bandwidth, both of which we assumed are limited resources (requirement **R4**). However, we can limit this problem with several methods. First, each on-going negotiation automatically means a risk of agreement. Given the decommitment problem, the agents should strive to find an appropriate balance between exploring the marketplace and getting into too many contracts. Second, we might also charge the agent for the used resources, so that agents would only open new negotiations when it really is necessary. Third, we might also have an absolute maximum number of concurrent negotiations that each agent is allowed to have at any given moment. The first option is embedded in the structure of the problem. We will also use the other two in some of our experiments.

Now, concurrent negotiation seems a possible approach, but to prove that it is truly useful, we have to explore how it compares to the group of mechanisms that are very popular in electronic environment, namely auctions.

2.3.1.4 Concurrent Bilateral Negotiation versus Auctions

In the wake of success of many auction web sites, such as `eBay.com`, online auctions and exchanges are seen as an essential part of e-commerce (Kambil and van Heck 2002). Annual online auction sales exceed already a few years ago $30 billion (David et al. 2005). This popularity makes auctions a baseline in all types of electronic markets and any research in a new or less popular mechanism, like concurrent bilateral negotiation, must justify itself by showing that the mechanism in question can provide something that auctions can not. This is not necessarily easy, since the popularity of auctions is not a coincidence: Auctions do have many desirable properties. In particular:

- *Legitimacy*: Most auctions have very simple and clear rules. This creates two major benefits:

[42] With leveled-commitment contracts (see section 2.2.3), there may be a cost involved in the latter action.

  - *Publicity*: During an auction, the bids can be open or at the very least the winning bid is usually public.
  - *Controllability*: Every participant can be certain that the rules are followed, since any deviation can be noticed. For example, the publication of the winning bid offers a way for the other participants to control that they were genuinely outbid.
  - *Neutrality*: In most types of auction, the identity of bidder is not relevant, but only the bid.

- *Analysability*: The simple and clear rules also make it possible to analyse auctions mathematically. To this end, auction theory (see Krishna (2002) for an introduction) is one of the most active research areas in game theory.
- *Efficiency*: Auctions usually allocate the sold item or service to the person or group that values it most. In economics, it is usually assumed that the person with the highest valuation knows the best use for the service and therefore auctions promote efficient use of resources.
- *Maximises profit of the auctioneer*: When used properly, auctions can provide the highest possible price for the seller or the lowest possible price for the buyer.

However, auctions are not a panacea. In particular, in dynamic markets such as ours, at least in some circumstances, the following properties might be considered drawbacks:

- *Inflexibility*: Clear and simple rules also mean that there are no exceptions or room for changes.
  - *changes difficult*: Once the auction has started, it is usually very difficult to change the rules of the auction[43] or the details of the auctioned task.[44]

[43] An auction usually has either a fixed deadline or it closes after a certain period of inactivity (i.e. when nobody increases the standing offer for a specified time). In a dynamic environment, where requirements can change at any time, both are problematic. In the case of a deadline, there can be no agreement before it. The bidder may be bound to his offer for a long time, if the deadline is distant and he is not outbid. During this time his requirements may change. On the other hand, if only a period of inactivity closes the auction, the exact end time may be difficult to estimate (Nguyen and Jennings 2005). Such open-endedness is very problematic for an auctioneer with a fixed deadline.

[44] For example, in eBay (`http://www.ebay.com`) a revision of the listing is only possible if an item has received no bids and the deadline is at least 12 hours away

  - *formulaic information dissemination*: When buying complicated and configurable tasks, the consumer must explain his requirements (winner determination rule) in detail to all potential providers, so that they can make meaningful bids. In some cases, companies or people may not want everybody to know what they need (Bajari et al. 2008) and so the only way to limit the flow of information, is to limit the number of bidders. Also the winning bid must usually be revealed, so that other bidders can ensure that they lost fairly. This may give useful information to the winner's competitors.
  - *one-way information flow*: Once the auction starts, the only information flowing between the parties is the bids. In some cases, the auctioneer might notice some error in his call for bids or his prefences when bids start. At that point, changes are no longer possible.
  - *heavy planning phase*: Since the auction itself is simple and clear, it often means that a lot of effort must be invested in the preceding phases (Bajari et al. 2008). For example, in a reverse auction, the consumer must find out what the relevant options and their characteristics are and what his preferences are, so that the winner determination rule can be written.

- *Neutrality*: In some cases, it might matter to the consumer, who the provider is and he may not want to make his preferences public.[45]

- *Complexity*: Although auctions are usually simple, they can become quite complex, especially computationally. In combinatorial auctions, for example, where bids on multiple services are allowed, the winner determination problem is NP-hard (for an introduction, see Lehmann et al. (2006)).

Concurrent bilateral negotiation, on the other hand, is very flexible, allows two-sided information exchange and also allows the party to control information dissemination and complexity and freely discriminate between different opponents. In addition, it can offer some strategic advantages. We will now discuss each of these benefits, one at a time.

(http://pages.ebay.com/help/sell/edit_listing.html). Also cancelling the entire auction (for any reason) becomes impossible when there is an acceptable offer (i.e. one that exceeds the reserve price) and there is less than 12 hours until the deadline.

[45]Neutrality was also mentioned as a benefit of auction mechanisms. Whether it is a benefit or a drawback, depends on the circumstances.

First, *flexibility* means that in concurrent bilateral negotiation, the agent can change its goals easily at any time. Since the deadlines and other negotiation parameters are private (requirement **R6**), changing them is always possible. Moreover, the agent always has several options to choose from (the opponents' last offers), if it suddenly needs to close the deal (Nguyen and Jennings 2005). In short, the agent can take into account any new information or change in environment or requirements very quickly.

Second, the possibility of *two-sided information exchange* has been recognised as one of the main advantages of concurrent bilateral negotiation (Nguyen and Jennings 2005; Rahwan et al. 2002). Another related advantage is the possibility of *controlled information dissemination*. Thus, the agent can start from a very general task description and reveal the details gradually. In addition, he can handpick the opponents to whom he gives any strategic information. The agent can also learn something from the opponents' offers that enable it to re-evaluate its preferences. For example, the option that the consumer thought to be the best, can prove to be too expensive or an option that it earlier hardly considered can prove to be very interesting. The flexibility allows the agent to 'change its mind', since it is not bound by any call for proposals. This is particularly useful when the agent does not know what the relevant options are or how much they cost.

Third, concurrent bilateral negotiation offers a way to *manage complexity in interconnected negotiations*. In combinatorial auctions, all parties first send bids to the auctioneer, who then tries to find the optimal winners among these bids. In contrast, in concurrent bilateral negotiation each party manages its own dependencies. This alone reduces a big problem into several small ones and effectively distributes the problem. In addition, each player can decide how complicated the dependencies he wants to consider should be and other players' dependencies do not have any effect on his problem.[46] This could be a major advantage for concurrent bilateral negotiation, which is why we will investigate also interconnected negotiations. However, the concurrent bilateral negotiation is no silver bullet: it does not necessarily offer an optimal solution, because suboptimal contracts between parties are not only possible, but also likely. This means that there is no guarantee that each service or service combination will go to the consumer who values them most. Also negotiation strategies and pure luck have an impact on

[46] Moreover, there is no deadline for bids, but the process is continuous, new parties can enter and old ones can leave at any time (requirement **R5**). However, there is a specific kind of auction called a continuous double auction that is able to manage this type of dynamism (Krishna 2002). Variations of this auction are used, for example, at all the major stock exchanges (Das et al. 2001).

the results. Therefore, concurrent bilateral negotiation will offer 'only' a solution that will produce reasonable results and will scale well to markets of any size.

Finally, the *strategic advantages* in concurrent bilateral negotiation come from two directions: a possibility to use different strategies with different opponents (Nguyen and Jennings 2005; Rahwan et al. 2002) and a possibility to use information from one negotiation in others (Nguyen and Jennings 2005). Since we assume that the providers are heterogeneous (requirement **R3**), it may well be reasonable to use different tactics against different providers. The consumer agent may, for example, adopt a stubborn stance against a low quality or unreliable provider, but be more willing to compromise with a high quality or especially reliable provider. And since the negotiation threads are separate (in the sense that the opponents do not know what is happening in the other negotiations), the opponents do not have to know that we consider them to be unreliable or to offer low quality. On the other hand, we can use information from the other negotiations. Thus, a good offer in one thread can be taken into account in others. Moreover, in a multi-issue negotiation, some new promising combination of properties that the agent did not think of, but an opponent did, can also be used in other negotiations.

Given these facts, it is hardly surprising that in practice private entities prefer negotiation even in situations where the law requires the public entities to organise a competitive bidding process (Rothkopf and Harstad 1994).[47] Since the private (unlike public) entities are free to choose, this result would indicate that negotiations are a more efficient means of transferring complicated assets than auctions and produce no worse results for the sellers (Rothkopf and Harstad 1994). Therefore we can be confident that our approach is a valid one. We will then discuss the details of this approach.

2.3.2 Managing Dynamism in Bilateral Negotiations

We start our discussion about concurrent bilateral negotiation strategies by investigating the simplest basic component of concurrent bilateral negotiation, the

[47]For example, Bajari et al. (2008) discovered that in California only about 15 % of the private building projects were awarded after a competitive bidding process. This was in clear contrast to public building projects where Californian law (like the law in many parts of Western world) requires a competitive bidding process for all significant projects and an auction was organised for almost all building projects. Public procurement processes also have other goals than efficiency (such as transparency and equality), so this is probably not as surprising as it may sound at first.

bilateral negotiation itself. We discussed the basic approaches to bilateral negotiation earlier (section 2.2.1). As we mentioned then, these basic approaches do not explicitly (or often even implicitly) take into consideration the changes in the environment. In game-theoretic and other mathematical approaches the reason is simple: the possibility of environment changes makes the model more complicated and the possibility of unforeseeable or unexpected changes is very difficult to model mathematically. However, as dynamic environments are seen as ever more important, a number of approaches that allow the agents to adapt to changes have been suggested. We will first discuss heuristic approaches (section 2.3.2.1) and then machine-learning approaches (section 2.3.2.2).

2.3.2.1 Heuristic Approaches

As in static negotiations, *the heuristic approaches* in dynamic environments try to translate the complicated reality into relatively simple decision rules that can be used to achieve good (although probably suboptimal) solutions. Here we will discuss one example, a market-driven agent (Sim (2002) and later versions of the model). This is a very interesting approach to negotiation in dynamic environments, because, as the name suggests, it takes the market situation explicitly into consideration when it makes decisions on concessions during the negotiation. When the market situation is tough (a lot of competition, a few opponents, deadline nearby or a strong need for the service), the concessions are bigger and when the market is more favourable, the concessions are smaller. In addition, market-driven agents consider their reservation prices as being flexible. Thus they are willing to pay more (accept less) if the market situation is hard and expect to pay less (get higher prices) when the market situation is better. We now discuss these two approaches in turn.

First, the concessions the market-driven agents make depend on their view of the market situation. Such an agent will take into consideration four different factors:

- trading opportunity,
- competition,
- deadline and
- its own eagerness to get the contract.

In more detail, *trading opportunity* is measured with two factors: the number of trading partners and the differences in utilities between the parties' last offers. Simple heuristics are used to estimate the probability that an agent will obtain a certain utility with at least one of its trading partners. They will then try to get the best possible utility while maintaining a reasonable probability of actually reaching it. Here *competition* is measured as a probability that the agent is ranked as the most preferred trading partner by at least one of the opponent agents. Again some simple heuristics are used in the estimation (the probability of being the most preferred partner is simply $1 - \frac{m}{m+1}$, where m is the number of competitors). Basically this means that a market-driven agent makes compromises according to the buyer-seller ratio in the market. For the *deadline* and *eagerness*, a simple time-dependent strategy is used. The eagerness factor ε determines how the concessions are made.[48] Here only two strategies are used: linear ($\varepsilon = 1$) and conservative (or boulware, $0 < \varepsilon < 1$), since the conceder strategy was likely to achieve lower utilities (although the lower risk of losing deals was noticed).

Second, market-driven agents in their newer versions are allowed to change their expectations of the outcome. Specifically, an enchanced market-driven agent (EMDA) (Sim and Wang 2004) can decrease its expectations in very tough market situations. To do this, it uses a fuzzy decision controller, which considers the factors above to guide decision-making. Yet another improvement is to allow the agent to increase its expectations when the market situation seems very favourable (EMDA2) (Sim 2004). To do this, it uses two additional fuzzy decision controllers that allow decision-making about whether or not to postpone reaching an agreement and, if so, for how long. The first decision is made by considering both the eagerness and the competition values; the higher these factors are, the higher is the number of good offers required to postpone the acceptance. The second decision uses the deadline and the number of opportunities; the further away is the deadline and the higher the number of opponents, the longer the agent is allowed to wait.

This approach is clearly very interesting in our context because the market-driven agents are very flexible and can adapt to many relevant changes in circumstances. However, it is relatively complicated and it is difficult to see, from this structure, if its adaptations are good in all relevant circumstances. Thus, it uses very simple heuristics and combines them into something that is no longer that simple to follow, because the different effects are all considered basically at the same time. We prefer a solution where different decisions are taken clearly in different times

[48] Basically ε is β in the descriptions of time-dependent strategies in section 2.2.1.2.

and/or by different components, and where the different factors are combined using a much more explicit analysis. This allows us to consider also cases that are not as clear cut as the cases Sim discusses.

Also the market-driven agents do not consider the opponent's negotiation tactics or his parameter values (like reservation price or deadline) and we have a market full of diffent kinds of sellers (requirement **R3**). Moreover, the approach does not take into account the possibility that the buyer agent's circumstances may change so that it does not need the service any more and should therefore be careful about entering into contracts (requirement **R2**). In addition, many of the details of the market-driven agents seem a bit ad hoc and no reasonable and consistent theory is offered to explain why these particular heuristics were chosen.

We do not find heuristic approaches satisfactory for our purposes, because the higher levels of our model will need relatively accurate information about the possible outcomes and success probabilities in all situations and heuristic models such as Sim's will not be able to provide such information.

2.3.2.2 Machine Learning Approaches

Since basically everything can change at any time and analysing the situation fully is simply impossible, we cannot find a perfect tactic that would work well in all circumstances. In addition, the situation may be so problematic that even devising successful heuristics at design time may be difficult. Therefore learning how to behave in different negotiation and market situations might be useful (Zeng and Sycara 1998). In addition, our environment, with potentially hundreds of agents, is probably too large, dynamic and unpredictable to devise heuristics that would work in every possible contingency. Thus the one possible way to cope with this is to allow the individual agents to improve their own performance (Sen and Weiss 1999).

In more detail, the literature on learning in negotiations has concentrated on two main things:

a. recognising the negotiation tactic the opponent uses and

b. finding a good tactic against it.

These problems are obviously connected. In competitive bilateral negotiations, especially in dynamic environments, the opponent can use very different tactics and there is no single best one that does well against all possibilities. Thus, the best tactic usually depends on the opponent's choice of tactic. This means that tactic recognition would seem to be the key. We will assume that once the opponent's tactic is known, we know an adequate tactic against it. In other words, in the following, we will concentrate on the first problem only.

So, as just explained, in *tactic recognition* the goal is to recognise the bargaining tactic the opponent uses. This can be an easy or very difficult task depending on what information we have available and what tactics the opponents are allowed to use. In particular:

- The quality and amount of available information is essential. Obviously in competitive settings the opponents do not explicitly explain what they are doing, so their tactic must be deduced from their behaviour (that is from their previous offers and reactions to our offers (accept, reject or make a counteroffer)). In addition, the environmental variables describing the market situation may be used (if available). However, some of the factors that may affect the opponent's tactic may not be available (for example his stock situation or the number of negotiations it is currently engaged in). If these unknown factors dominate the tactic, it may well be impossible to guess what happens next. On the other hand, if the most important factors are available to all, tactic recognition *may* be possible.

- If the opponents are allowed to use only a small number of tactics, which are easily distinguishable (e.g. if there are only two possible time-dependent tactics with $\beta = 10$ and $\beta = 0.1$), tactic recognition is trivial. On the other hand, in completely open settings (any tactic allowed) the tactic cannot probably be estimated to a comfortable certainty, before the negotiation is over, if even then. In this latter case, we might be able to produce an educated guess as to what will happen next given the history of negotiation, identity of the opponent, other information we have, and our earlier experiences, but we can never be sure, since the opponent can change its mind and tactic at any given moment or use a tactic we have never seen before.

In this work, we are somewhere in between the two extremes in terms of both of these factors. Our consumer agents do not have a lot of information (requirement **R6**). We assume that parties do not know each other's deadline but that they may know other parameters (for example quality and reservation price).

Now, there are two broad types of machine learning techniques that can be used to recognise an opponent's tactic: classification and regression. In *classification* the opponent's tactics are classified to one of a finite number of categories given their behaviour so far (and other relevant factors). The opponents in one group can then also be assumed to behave in a similar manner in the future, so that it might be possible to learn one tactic that works well against all of them. Obviously this approach is not optimal in all situations, but it should provide us with some idea about what will happen next. In the literature, classification has not been used to classify opponent tactics, but it has been used to estimate the opponent's utility function in multi-issue negotiations (Bui et al. 1999; Chajewska et al. 2001; Coehoorn and Jennings 2004)[49] or the reservation price of the opponent (Zeng and Sycara 1998).

In *regression*, the goal is to estimate a value of some (continuous) variable or function (for example, the next offer, deadline or reservation price), given the events so far. It is usually assumed that once the regression yields good estimates, a good countertactic is also known. Unlike classification, regression has been used in tactic recognition in negotiation contexts. In particular, Mok and Sundarraj (2005) use regression to estimate the parameters of the time-dependent heuristic tactic the opponent is assumed to be using. After a reasonable degree of accuracy has been achieved, the optimal tactic against the estimated tactic is used. Hou (2004) also uses non-linear regression on the opponent's previous offers to estimate the opponent's deadline and reservation price, and then chooses an appropriate countertactic. Hou assumes that the opponent is using one of the heuristic tactics discussed in section 2.2.1.2, but he restricts quite drastically the number of possible parameter values (for example there are only four different deadlines and four possible reservation prices). The agent tries to first classify the tactic the opponent is using and when it has a good idea of that, it tries to use regression to estimate the parameters.

However, in any realistic setting, the providers would be using a very wide range of negotiation tactics, some of which would be new or very rare and some of which would use information that is not available to us at all. Some of these tactics might also be very difficult to distinguish from some others. And in a completely open

[49] Bui et al. (1999) work in a cooperative environment and classification is used to predict other agents' preferences in an effort to reduce communication. Chajewska et al. (2001) try to elicit the opponent's utility function from the observed negotiation. The setting is competitive, so the estimates are used to increase the learning agent's own utility. They have existing partial utility functions that are used as classes. In a similar vein, Coehoorn and Jennings (2004) use kernel density estimation to estimate the utility function of the opponent to make efficient multi-issue negotiation trade-offs.

environment, there is also nothing to stop our opponent from changing its tactic completely in the middle of the negotiation or employing random noise or any number of other methods to make tactic recognition an impossible task. Although we might be able to extend the work of Mok and Sundarraj and get reasonably good results in our restricted environment, such 'perfect' tactic recognition is not attainable in the general case.

Therefore, we contend that machine learning techniques would be very problematic in even a remotely realistic setting. We think that a more reasonable 'countertactic' in such environments would be based on estimates of the opponent's reservation price and other parameters, not on the offers the opponent makes, because single offers may not relay that much useful information in practice, only the fundamentals (the parameters) behind the tactics matter. Also, if we have enough opponents to choose from, we may not have to succeed in each and every negotiation but we may, for example, employ tactics that work against some opponents and fail miserably with others.

However, we do need some information to use any sophisticated strategies in the higher levels of our models. One good approach might be to set an offer to a certain level (based on actual values or estimates of quality or other attributes) and then use empirical data to estimate success probabilities in a given negotiation. In the perfect world, the buyer agent would be able to recognise the tactic the seller uses or even know it in advance (because of many earlier encounters). It would be interesting to see how well our approach will work in such cases and then we can remove or limit this information and see what effect that has.

So, in our setting, the sellers will have a small number of heuristic tactics and we have a countertactic for each of them. This countertactic is based on the characteristics of the seller (quality), the general structure of the tactic and maybe used against many types of tactic. However, the tactics will use randomness and/or behavioural aspects so that the regression models we just discussed are not applicable. We also will not make any use of any classification schemes (although they might sometimes be useful). Instead, we assume that sometimes the buyer might just know the tactic the opponent uses in advance (no learning takes place during the negotiation) or at least it knows the frequencies and types of negotiation tactics. This can be at least in some circumstances realistic. We hope to show that although information about opponent tactics may indeed be very useful, we can get quite good results even without such information. Of course our approach still requires that we have an idea what sort of tactics there are in the market and their probabilities.[50] But we believe, this approach might well offer a reasonable approach to bilateral negotiations in open environments, although admittedly, we will be using it in a rather restricted setting.

2.3.3 Concurrent Bilateral Negotiation

Our buyer agent can be engaged in several bilateral negotiations in parallel in order to secure a good deal in dynamic markets. In this section, we discuss the current literature on such negotiations. However, before we discuss any specific issues, we briefly describe the three general models that have been introduced in the literature (section 2.3.3.1). After that we will discuss the literature on three questions:

- *opponent selection and concurrency control*: How many opponents to negotiate with? How to choose the opponents for these negotiations? (section 2.3.3.2)
- *relationship between negotiation threads*: Who makes decisions on accepting an offer? (section 2.3.3.3).
- *negotiation strategies*: What offers to make in each negotiation and at what level is this decision made? (section 2.3.3.4). The first and third questions are obviously linked. The strategies we use clearly affect the way we should choose the opponents and vice versa.

2.3.3.1 The Basic Models

We start by introducing the three models of concurrent bilateral negotiation discussed in the multi-agent system literature. They are similar in many respects, but there are also some significant differences. Some are also more explicit on the details we are interested in than others.

In particular, Rahwan et al. (2002) were the first to propose a concurrent bilateral negotiation model. They have a number of bilateral negotiating threads (called sub-negotiators) controlled by one coordinating agent. In this model, each

[50] As we will discuss in more detail in section 9.1.2, the countertactics we use do not care what offers the opponent makes during the negotiation but instead they use a distribution for the lowest offer the opponent is going to make and set their target price accordingly. With many negotiations, getting this absolutely right is not essential.

sub-negotiator conducts a one-to-one negotiation with a different opponent. The coordinating agent coordinates the efforts of these different threads. To facilitate this, the sub-negotiators report to the coordinating agent after receiving a response (offer, rejection or acceptance) from their opponent. The coordinating agent then evaluates the situation and issues instructions to sub-negotiators accordingly. The most important of these instructions is either an order to continue or stop negotiating, but they can also include a change in the negotiation parameters. However, Rahwan et al. do not develop this idea very far and the strategies their coordinating agents use are very simple and somewhat unrealistic.[51] Also these strategies assume a quite one-sided protocol: the offers made during the negotiation are binding for the sellers, but the buyer is allowed to decommit at any time. This lack of balance is not discussed in the paper.

Building on this work, in a series of papers Nguyen and Jennings (2003, 2004, 2005) propose a similar model for the buyer agent. They use time-dependent heuristic strategies (boulware, linear or conceder) in their bilateral negotiations. Their model is also developed further. In more detail, the architecture consists of three main components: a coordinator, a number of negotiation threads and a commitment manager (see figure 2.1). The roles of the coordinator and negotiation threads are similar to Rahwan et al, although the interplay between them is defined in more detail and the coordinator is also given a right to decide the strategy the negotiation thread uses. The big difference is the commitment manager, which makes all centralised decisions on commitment or decommitment for all the threads (Nguyen and Jennings 2005).[52] This became necessary when the bias for the buyer was removed and both parties were given an equal right to decommit. The commitment manager, in close cooperation with the coordinator, approves any acceptance of an offer (ensuring that only one offer gets accepted at a time) and it also makes the decisions on decommitting in case more than one of the offers gets accepted by the opponents or when a significantly better contract has been found. We will discuss this in more detail later.

Another relevant set of papers is written by Li et al. (2004, 2005, 2006). Their approach is somewhat different. They do not have an explicit coordinator module at all, but instead coordination occurs by calculating the expected value of an outside option for each thread. An outside option is a game theoretic concept used in bargaining models. It means simply the best available outcome outside

[51] We will return to these strategies in section 2.3.3.4.

[52] The commitment manager was introduced in Nguyen and Jennings (2005). The earlier versions of the model were one-sided in similar way to Rahwan et al. (2002) (i.e. the buyer can decommit, but the seller cannot).

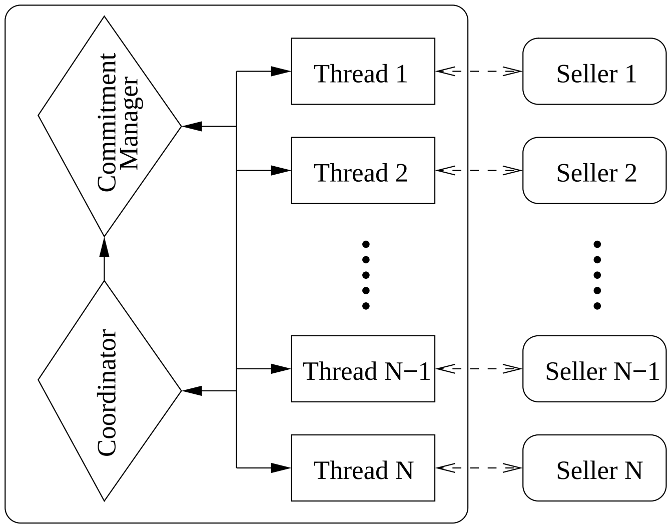

FIGURE 2.1: System Architecture in Nguyen & Jennings (2005).

of the current negotiation (should that fail). In concurrent bilateral negotiation, an outside option is the highest expected utility of all other negotiation threads and the possible later negotiations. Since in negotiation the outcome is always uncertain, Li et al. describe several methods for making this estimate. They are less interested in the details of the coordination and it is, therefore, unclear how this would deal with a situation in which more than one of their latest offers gets accepted by the different opponents. Again, we discuss these issues later in this section.

Now, in general, the model introduced by Rahwan et al. and further developed by Nguyen and Jennings seems to offer the best starting point for our work. A controller needs to manage a group of negotiators, handle accepting offers and making decommitments and the negotiators should report about their progress every turn. This is very good and it can be easily extended to cover other problems. The work by Li et al. is also interesting, however, and, as we see later, some of it does find its way to our model. It should also be noted that one-to-many negotiation can easily be changed into many-to-many negotiation by allowing both parties to use one-to-many negotiations (see figure 2.2). However, we will stay in the one-to-many model, since there are many inadequately explored issues in the one-to-many situation and we wish to concentrate on those.

Having introduced the basic models in concurrent negotiation, we will now turn to three broad problem areas that, on one hand, are essential to concurrent negotiation and, on the other hand, are unsatisfactorily dealt with in the current literature. We start from the opponent selection and concurrency control.

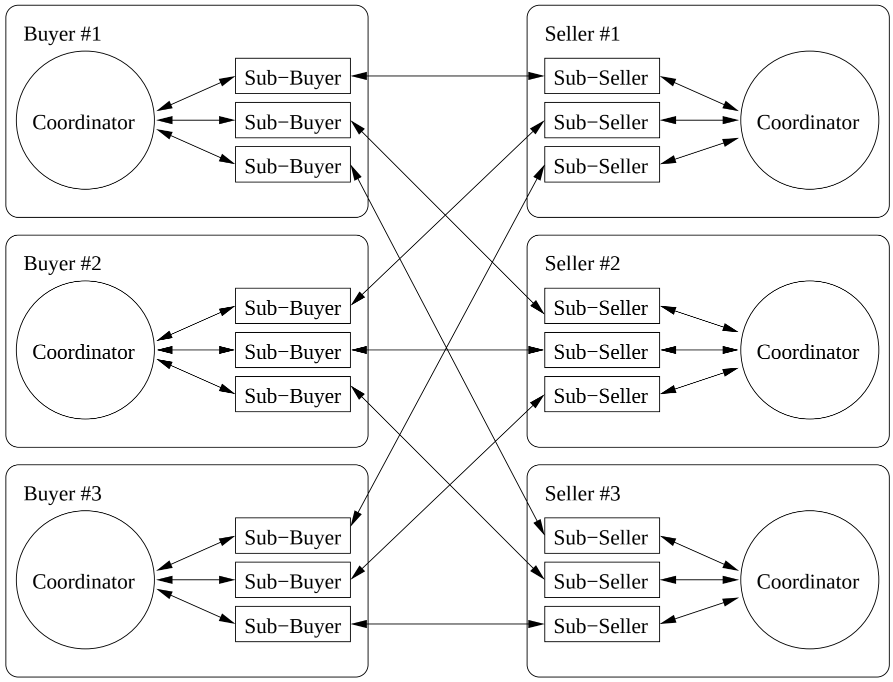

FIGURE 2.2: A many-to-many market in Rahwan et al.(2002).

2.3.3.2 Opponent Selection and Concurrency Control

Concurrency control means selecting the number of negotiations to engage in concurrently and *opponent selection* refers to choosing opponents for each of these negotiations. All the papers discussed in this section assume that the agent can negotiate with all possible providers. Although this may be possible in some situations, we do not think that this is generally the case and we assumed that resources to be used in negotiation are limited (requirement **R4**). On the other hand, our sellers are going to be heterogeneous (requirement **R3**), as probably are the sellers in most markets. For these reasons, it is therefore necessary to choose the number of opponents and the actual opponents to negotiate with. Since these issues have not been addressed in the concurrent bilateral negotiation literature, we have to look elsewhere.

Unfortunately we could not find any literature on concurrency control, nor could we specify any other problem, which would have been discussed in the literature and would be similar enough to be useful for our purposes. However, as we explained, concurrency control is important. Also it matters because of a risk of getting into too many contracts. In some environments (cheap/free decommitment), it might not be a problem, but where the decommitting is very expensive, a useful buyer agent should be more careful. For this reason, we will develop methods for concurrency control in our work. Using decision theory seems a reasonable starting point for that work, because it allows us to consider different aspects of the situation in a unified way.

We could not locate any literature on opponent selection either, but we were able to find another problem that is close enough and has plenty of literature on it. This other problem is service selection. In the service selection problem, an agent has a certain task to be performed and has to choose services that satisfy its needs optimally. However, there are two differences between opponent and service selection:

- *Dependencies*: Instead of finding an optimal service to fulfil a certain task, the service selection task is usually performed over several interconnected tasks and the goal is to find the services that are interoperable and produce the best possible result. Since we assume that our negotiations can have this type of connections, this might actually be quite useful, although at this level we are only interested in selecting the best service for the one task at hand. We will return to the interconnectedness later. However, since it is the dependencies between services that make the service selection an interesting problem, the actual selection of one service usually plays a very minor part in these papers.
- *No negotiation*: The work of service selection usually assumes that there is no negotiation on the price or quality, but that these have been set by the providers on a take-it-or-leave-it basis. Therefore the price is usually known in advance exactly and quality is an estimate (because it can vary). However, for us it is other way around: each provider has a fixed quality (but in some case may decommit and therefore not perform at all) but the price is a result of negotiation and therefore, it is uncertain.

With these differences in mind, we now discuss the literature. Service selection is usually based on multiple criteria, so the literature uses many standard techniques for such situations, such as multi-attribute utility theory (Keeney and Raiffa 1976) and decision theory (Raiffa 1987). Also some other approaches have been used. We will now discuss how these could be applied to our problem.

The basic idea of *multi-attribute utility theory* is that the overall value of a service consists of the values of its performance on the relevant characteristics, for example its quality and reliability. Usually, simple additivity is assumed, so that the value of a service is equal to:

$$v(s) = \sum_{i=1}^{k} w_i v_i(s),$$

where $v_i(s)$ is a normalised value of the characteristic i for service s, k the number of characteristics, and w_i is the relative importance of that characteristic for the decision-maker. Normalisation means that given any possible value of characteristic i the relevant value function V_i has a value in the interval $[0, 1]$. The better the attribute, the higher the value; however this relation does not need to be linear but can take any form. Another common assumption is that $\sum_{i=1}^{k} w_i = 1$. Now, it is clear that the critical task in multi-attribute utility theory is to find the appropriate weights for the overall value function. There are many standard techniques for deriving these, probably the best-known being the analytic hierarchy process (Saaty 1980). The service selection papers that use some version of multi-attribute utility theory include Seo et al. (2004) and Zeng et al. (2003).

However, in situations where the utility is quite clearly not a sum of characteristics, but something else, multi-attribute utility theory does not work too well. Here, we defined the consumer's utility as: $U(s) = V(q_s) - p$, where $V(q_s)$ is the value of service s and p the price (one of the characteristics). Now, if we also consider the provider's reliability w_s, the probability that it will provide the agreed service at the agreed time, we get:

$$U_{consumer}(s) = w_s[v(q_s) - p_c] + (1 - w_s)fee,$$

where fee is the decommitment fee the provider has to pay if he fails to deliver the service. Now, the different characteristics are no longer additive, so the basic form of multiattribute utility theory is not applicable.

In contrast, in the *decision theory* approach the decision-maker will simply use his utility function to calculate his expected utility for using each possible service and then he just chooses the option that provides the highest expected utility (Collins et al. 2001).[53] Here, the expected utility would be calculated using the equation for utility and replacing w_s, q_s and p_c with the values from the service in question.

[53] The expected utility is of course very strictly confined to the situation at hand (as discussed in section 2.2.1.1) and since the situation of Collins et al. is significantly different (they have a reverse auction and multiple interconnected services), it is not useful to go through their model in detail. Instead, we will discuss some problems of using this approach in our environment.

However, although each provider's quality of service q_s and reliability w_s are assumed to be known (requirement **R6**) and also the decommitment fee fee and the consumer's value function $v(..)$ can be assumed to be known by the consumer agent, the problem is the price p_c, which is *not* usually known at the time of opponent selection by anyone, since it will be determined by a negotiation *after* the opponents have been selected. The provider-specific price distributions are also likely to be very hard to learn even over considerable time, since they depend on the negotiation tactics used by the parties and their negotiation parameters (such as deadlines and reservation prices), which are likely to vary between negotiation encounters. However, we can make this choice properly when we use negotiation tactics that make the same offer again and again. Then we know exactly what the price is going to be if we are successful in the negotiation and a contract is formed. Of course we need to consider and be able to estimate the probability of being successful in a negotiation, z, so we get the following:

$$U_{consumer}(s) = z(w_s[v(q_s) - p_c] + (1 - w_s)fee).$$

In order to get a useful estimate to z, we would need to know what negotiation tactic the opponent is using or at least we need to know the possibilities and the probability distribution over them. If we have very little or no information about the opponent tactics, the meaning of this extension will be very small. If z is roughly the same for the all the opponents, it will make a very small impact on the selection. Despite these problems, the decision theoretic approach seems applicable to our problem and so we will adopt it (see section 10.1.2.1 for details).[54] We will of course use also other (much simpler) opponent selection techniques, but we use the decision theoretic approach as a pinnacle of opponent selection, the most sophisticated technique.

Using a decision theoretic approach in both opponent selection and concurrency control means that we will need a significant amount of information about the market and the sellers. Our information requirements (requirement **R6**) have made it possible for us to consider and choose decision theoretic approaches to some issues, although such information may not always be readily available. We,

[54] There are also other approaches to service selection, such as the preference ranking organisation method for enrichment evaluations (PROMETHEE) (Brans and Vincke 1985), which was used for service selection by Seo et al. (2005), and constraint satisfaction/optimisation (used for service selection in Lin et al. (2005)). They all suffer from the same problem: the best goal function depends on p_c, which may be difficult to estimate in advance and they are not directly applicable to our problem. The PROMETHEE is quite complicated and there are very few constraints in our setting.

however, think that it is good approach to see first what works when there is good information and then make the setting more challenging (and realistic) by taking some of that information away and trying to cope with less information. We also think that although our information needs are significant, they can, with some additional work, often be relaxed with a limited loss of performance. We also use less sophisticated methods in some cases to see how they perform. Of course this work should be continued in future work to settings with less and less information.[55]

2.3.3.3 Reaching Agreements

The *relationship between the various concurrent negotiation threads* is a key issue in concurrent bilateral negotiation. To this end, we discuss two points:

- how is the decision to accept the opponent's offer made and
- what happens if one or more of the offers gets accepted?

Now, the solutions to both of these questions vary in the literature. In particular, Rahwan et al. (2002) use autonomous negotiation threads that each try to reach an acceptable agreement on their own. It is, therefore, the thread itself that decides when to accept the offer and when to continue negotiation. The coordinator can order threads to quit negotiating, but not to accept or reject a particular offer. The coordination between different threads is achieved ex post (after agreement), as the coordinator uses one of four strategies to select the agreement to be bound by. Strategies are:

- *desperate*: as soon as an acceptable offer exists, it is accepted and all other negotiations are ended. If there is more than one acceptable offer, the one with the highest utility is selected.
- *patient*: threads that have an acceptable deal are asked to wait while other threads finish their negotiation. Then the result offering the best utility is selected and others are decommitted or ended.[56]

[55] We will discuss this in more detail in section 12.2.2.

[56] The strategy guarantees the best possible deal, but does not care about time constraints. One variation is that at the deadline the best offer so far is selected and the others are rejected.

- *optimised patient*: as patient, but the outcomes so far affect the reservation prices of the remaining negotiations, so that each new agreement is better than earlier ones. Thus the utility constraint is updated to the highest offer so far.
- *strategy manipulation*: the coordinating agent may modify the negotiation strategies of different sub-negotiators at run time.

All but the last one are very one-sided, since they assume that the providers obediently wait until the consumer makes a decision on which offer to accept and that negotiator can decommit from any other agreements without any problem at any time. It seems that the agent does not need to pay anything to decommit and the possibility that the opponent accepts the agent's offer is not discussed in the paper. The most interesting approach for our purposes (strategy manipulation) is not explained in any detail.[57] In the model it has no strategic relevance whether the consumer agent's offer was accepted or whether it was the consumer agent that accepted the provider's offer, since in either case coordination is done ex post.

A different approach is offered by Li et al. (2004), in which all negotiation decisions are made in the negotiation threads, but the coordinator does influence the tactics used in each thread already during the negotiation. The coordinator estimates the outcomes of different threads and uses these estimates to set the reservation utility for each thread.[58] In order to make the estimates several big assumptions are made: the buyer agent is assumed to know the distribution of the providers' reservation prices, the newcomers' arrival probability and their items' value distribution. As Li et al. have done, we assumed that the sellers can exit and more sellers can enter at any time (requirement **R5**). Any reasonable negotiation model that works in the market environment should usually consider this requirement and consider any good contracts the latecomers might bring with them. It may not be a good idea to take a mediocre deal if there are good providers probably coming later. Li et al. use analytic approach and knowledge of many distributions to calculate their estimate, but we will instead use empirical data collected on previous runs,

[57]Only one example is given and that is also very one-sided: Once an acceptable offer exists, all the other threads are ordered to send a take-it-or-leave-it offer with higher utility than the one already reached. If none of these offers gets accepted, then the original agreement is selected. However, if one of these new offers gets accepted, it will get selected and if more than one offer is accepted, one is selected at random.

[58]The reservation utility is equal to the expected utility of the outside option (highest expected utility of all other negotiation threads and negotiation threads starting in the future). They use a simple time-dependent heuristic strategy to find out what offers to make. They also offer different methods to estimate the outside options and compare the results.

because we believe that is more realistic. Moreover, it allows us to consider settings in which such distributions for future expected utilities might be difficult to derive.

A problem with the model by Li et al. is that they do not discuss the coordination issue any further, so either they assume that once one thread accepts an offer, others are automatically ordered to quit (and there is enough time to do this) or that the consumer agent is able to decommit from all extraneous contracts for free. Since cases where the consumer agent's offer is accepted and the consumer agent accepts the offer are not separated, the latter seems more probable. This is inappropriate for our settings because decommitment policies and how they should affect the parties behaviour is a central theme to our work and this particular problem (extra contracts) is an important part of what we investigate. Therefore, for us, it is obvious that simultaneous contracts are possible and the buyer should have to pay a decommitment fee for any extra contracts it does not need.

Finally, Nguyen and Jennings (2005) use very explicit coordination mechanisms and, unlike the other models, all decisions on acceptance are made by the coordinator. This coordinator gets all the opponent offers from the negotiation threads and then decides whether to accept the best of them or to continue negotiating. The negotiation thread's task is then to implement this decision. However, more difficult situations occur when more than one offer gets accepted by the provider agents. In such cases, the consumer agent just decommits from these extra contracts and it seems that it can do so by paying a decommitment fee.

We require that our model must have an unbiased protocol. Therefore, a situation where one party can and the other cannot decommit is not acceptable to us. On the other hand, we think that, if and when a contract has been made, the decommitment should usually cost something. Consequently, we think that the last approach by Nguyen and Jennings is the most appropriate starting point for our work. Their model is also the most explicit when it comes to the mechanics of committing and decommitting. The other models seem to have suprisingly little to say about these topics, although they seem essential to any concurrent bilateral negotiation model.

2.3.3.4 The Negotiation Strategy

A *Negotiation strategy* is a collection of negotiation tactics for on-going bilateral negotiations. We discussed some issues involved in devising such tactics for bilateral negotiation in sections 2.2.1 and 2.3.2. Here, however, we will discuss how having many concurrent negotiations may affect these tactics. Also here it seems that state of the art is not very advanced.

In Nguyen and Jennings, both parties to all negotiations use time-dependent heuristic negotiation tactics. The buyer agent can use three different strategies (conceder, linear and boulware) and there are two different types of sellers (one uses a conceder strategy and the other a boulware strategy). The buyer knows the share of these types in the market and he also has probabilities of successful negotiations and expected utilities for each combination (conceder/conceder, conceder/boulware, etc.). The buyer's negotiation thread then calculates the expected utility for each possible strategy and chooses the one that maximises it. In this case, strategy choices in one thread affect the choices in others, since the type probabilities are updated after a strategy is chosen for each thread using Bayes' rule (assuming that the opponent in the current negotiation thread is of the type that has higher probability in the selected strategy). The thread will also try to guess whether the opponent uses a conceder or a boulware strategy. Since this is quite an easy distinction to make, the buyer soon knows the type of each seller and can select the optimal strategy against him. This seems too simple for our use, although the basic idea of trying to recognise the opponent strategy and then choosing a good strategy against it, is obviously what we will plan to do in a bit more complex environment, where distinctions are more difficult to make. In addition, the role of the decommitment fees in deciding how many negotiations to use is not discussed.

In Li et al. the estimated outcome of other negotiations affects the negotiation's reservation utility. Since different negotiations have a different set of these other negotiations to consider, the reservation utilities can vary in different negotiations. In addition, the value of the opponent's item/service can be different. This means that the reservation prices are different in each negotiation. In particular, Li et al. use a basic time-dependent heuristic strategy with constant β. However, since reservation prices for each negotiation thread are calculated in every round, the reservation price can increase or decrease in each recalculation, which means that an offer can increase or decrease. The idea of adjusting the reservation utility based on the expected outcome in other negotiations is quite interesting. A nice property of this approach is that the strategy will be 'automatically' harder against low quality producers and softer against high quality producers.[59]

[59] As we learned in section 2.1, the consumer's utility is the difference between the value of a service and its price. With lower value (quality), also the price needs to be lower to achieve a certain reservation utility. This means that with the low quality providers, the reservation price will be lower.

However, we think there is a plenty of room for improvement here. We will use the idea of using expected outcomes in other negotiations as a basis for reservation price, but we do not use that directly with on-going negotiations but with future negotiations. So if the buyer expects to have good offers later, it will use that as the reservation utility it needs from negotiations it is currently engaged in. Also the idea of demanding the same (expected) utility from all opponents is going to be used in our most advanced concurrency strategy. In general, the papers seemed to have a bit of trouble deciding where the negotiation tactics should be set: at a controller or at a negotiator and what role each of them have in this process. For this thesis, this is also an issue that should be addressed. If the model is designed to be extendable, it needs to have a very clear view on what decision is made by what component and when. Moreover, these models does not address the situation where the buyer might change its mind about wanting a service at all (requirement **R2**). This is something we feel should be also incorporated into the buyer agent's model.

2.3.4 Interconnected Negotiations

In most of the literature, each automated negotiation (or at least a set of negotiations on a single service) is considered in isolation. There are some exceptions such as Zhang et al. (2005), but none of this work is done in the context of concurrent bilateral negotiation and within the framework of a concurrent bilateral negotiation model. We believe that managing interrelations on different services should be a significant part of any such model. On the other hand, our main interest will be on improving the management of concurrent negotiation on the same service and we will therefore not make hugely complicated models here. Instead, we will just discuss two basic cases of interconnectedness between services: The different services can be either substitutes or complements to each other (see section 10.1.2.1). The services A, B and C are substitutes when only one of them is needed and they are perfect substitutes if the decision-maker is indifferent between the services, in other words, it does not care if it gets a service A, B or C as long as it gets one. The same services A, B and C are complements when all three are needed for the highest utility. We will discuss these two types of interconnections in our work.

2.4 Summary

We have discussed the relevant literature and made our observations on what works in the state of the art and where there is room for improvement. We will now summarise our findings in two parts. First, we discuss the commitment models (section 2.4.1) and then the concurrent bilateral negotiation strategies (section 2.4.2).

2.4.1 Commitment Models

We have discussed the literature about bilateral negotiation and decommitments. On bilateral negotiation, we concluded that the heuristic negotiation tactics will be simple, light (limited negotiation resources, requirement **R4**) and easy to use and will allow us to consider all the relevant requirements. They allow us to have costly and time-consuming service preparation (requirement **R1**), because we can easily translate such limits to deadlines and reservation prices.[60].

And although we did not discuss it explicitly, our marketplace should also fulfil the requirements of openness (requirement **R5**), heterogeneity (requirement **R3**)[61] and incomplete information, which does include the availability of some basic information for all parties (requirement **R6**). Moreover, the marketplace should allow for changes in circumstances (requirement **R4**). These requirements are not especially relevant to our literature review, because most approaches we discussed could have filled these requirements.

We also concluded that the leveled-commitment contracts is the best approach to manage changing circumstances (requirement **R2**) and the necessary decommitments. We will investigate the effects of four different decisions that the parties make in our market environment (performance, reliance, contract and selection) and try to find out what effect the different decommiment policies (rules on decommitment fees) have on the outcome in terms of common good. For the most part, we use the decisions as they are described in the literature. The literature on many decisions is limited in the sense that models used (if any) are often quite simple or even non-existent and the discussion revolves around general principles. We will use these principles and basic settings and adapt them to our dynamic service market setting.

[60]We discuss this process in detail in section 3.2.

[61]Again, the latter can be achieved by choosing some parameters at random, see section 3.2 for details.

Because the papers are published in the law and economics literature, they often consider incomplete information as a question of evidence to be settled in a court of law and do not consider other options or more complicated settings. We will therefore extend these models to settings where information is incomplete and create decommitment policies that can address this situation quickly without a need for third parties or complicated procedures. Moreover, the discussion in the literature is often brief and limited to identifying the optimal policy in the circumstances. Sometimes, this policy may be complicated or depend on information that is not readily available and it is, therefore, useful to discuss also other policies and investigate how well simpler or otherwise sub-optimal policies might perform in different settings. We also investigate the role of some other characteristics and rules of the marketplace such as population sizes or possibility of re-entry, to name a few.

Our main focus will be on the performance decision, because it is the most established of the four decisions and because legal rules may give us pointers for deciding what to do in special circumstances or under incomplete information.

2.4.2 Concurrent Bilateral Negotiation Strategies

We have discussed the literature on concurrent bilateral negotiation and have given an example of how interconnected negotiations may be managed. We have identified several methods in the current state of art, which will be useful to us and there was, for example, models where the openness (requirement **R5**) was considered to a reasonable extent, although we considered having distributions of future opponents a bit excessive and will use empirical data instead. The existing literature also provided adequate support for our requirement of costly and time-consuming service preparation (requirement **R1**). Also, many of the models we discussed made quite strong assumptions on available information. We have made our own assumptions (requirement **R6**) that are, compared to the state of the art, quite strict in some places and quite generous in others. We believe, however, that at least some of our assumptions can be relaxed in future work with limited loss of performance.

However, we have also identified some quite significant gaps in the state of the art. These are mostly connected to the managing concurrent negotiations:

- concurrency control: there is no discussion of selecting the number of concurrent negotiations in the literature, although this seems quite essential task in a concurrent negotiation, because it directly influences the number of extra contracts the buyer agent may have to decommit from and also because the concurrent negotiation may be computation-intensive and resources that can be used for negotiation usually have limits (requirement **R4**),

- opponent selection: there is also no discussion of choosing the best opponents to have the concurrent negotiations with, although in most markets, the sellers are highly heterogeneous (requirement **R3**),
- negotiation tactics: although the literature had some good ideas, there is some room for improvement. We need to be able to estimate the outcomes for our negotiations and chances of success and none of the tactics in the literature offered us that. Also all the concurrent bilateral negotiation models used very simple negotiation tactics on the opponent side (often simple time-dependent tactics) and we want the sellers to use more diverse tactics.
- changing circumstances: a case where the buyer (and the seller) have a probability that they may not want the contract to go through in the end but want to decommit at some point (requirement **R2**), has not been considered, in the literature, especially not in the context of concurrent negotiation,
- interconnected negotiations: interconnected negotiations have not been considered in the context of concurrent bilateral negotiation

Now, to address these shortcomings we need to develop a concurrent bilateral negotiation model that includes support for interconnected negotiations. We also need to pay extra care to designing the structure and inner workings of a negotiation-coordinating controller: how these mostly new functions can be made to work well together, so that new strategies and methods can be easily added. We also need to work on the division of labour between the different levels of the model. Each such level should have their own areas of expertise and they should not usually second-guess each other's recommendations, except if one level has a wider view and needs to adjust a recommendation, but even then it should not just do whatever it pleases but use the expertise of the other levels.

Part I

Commitment Models

Chapter 3

The Marketplace Model

Now that we have discussed the state of the art and its shortcomings, it is time to explain in detail how we are going to address some of the issues we have identified. In other words, we will now move to discuss our research. As discussed in the introduction, we will explain our work in two major parts each consisting of several chapters. The common thread through all our work is to consider how the decommitment policies affect the way an intelligent market participant should behave in a marketplace. In part I, we are going to discuss system-level effects of decommitment policies. If and when the decommitment policies influence behaviour of the individuals in the market and if and when the market outcomes (and hence the welfare effects of a marketplace) follow from the actions of the participants, it can be said that the decommitment policies influence the common good the market can produce. In other words, in part I, we are interested in how the way the market participants adapt to the different decommitment policies affects the common good or the sum of utilities of all participants.[1]

Specifically, in part I, we use a relatively simple market setting and keep the adaptations the parties use simple and straight-forward and investigate how the decommitment policies affect the big picture to get a clearer view on the role the decommitments and decommitment policies can play in an electronic marketplace. In more detail, we discuss the underlying theory and our results in the market setting in five parts. First, in this chapter, we introduce our marketplace model and its implementation and in the four subsequent chapters we discuss each of the four major decisions that the parties make to adapt their behaviour to the

[1] We discuss the other aspect, the individual-level strategies mainly in the part II, although some basics are discussed also in this part.

decommitment policy in force (this corresponds to research contributions **C1**–**C4**):

- performance decision (chapter 4, contribution **C1**),
- reliance decision (chapter 5, contribution **C2**),
- contract decision (chapter 6, contribution **C3**), and
- selection decision (chapter 7, contribution **C4**).

For our marketplace model, we consider a market of buyers and sellers for one service. We refer with subscript b to a single buyer (consumer) and with subscript s to a single seller (provider) in this market. We are especially interested in their utility, U_b and U_s, respectively. The time t is discrete and divided into turns. We assume that all participants expect the delivery of the service to occur at the same time $t_{delivery}$ (there can be separate markets for different delivery times and other services but we are only interested in one market). We first explain how the market works (section 3.1), then we discuss how the negotiations proceed and how the parties get the parameters for the negotiations (section 3.2). We will then discuss how we make the parties consider issues associated with decommitment on the contracts that have been formed (section 3.3) and how we implemented this system (section 3.4).

3.1 Matching and Entries

The buyers and sellers in the market are paired at random by the marketplace. This means that each provider will be given one consumer to negotiate with (and vice versa). The pairs then negotiate for 100 turns on the price of the service. Both parties use simple exponential time-dependent heuristic tactics (section 2.2.1.2), in which the parameter β is selected at random. Once all negotiations finish, the parties remaining in the market are again matched at random. This process (from the random matching to the end of negotiations) is repeated 10 times. The entries and exits can occur at any time, but the parties are only matched at turns that can be divided by 100 without a remainder (i.e. $0, 100, 200, ...$). If there is an unequal number of buyers and sellers in the market, some members of the larger population will not get an opponent and will have to wait until the next matching. The contracts are performed when the negotiations end, thus $t_{delivery} = 1000$.

Size	Initial Size (n_0)	Entry Intensity (i_b)
Tiny	0	0.02
Very Small	0	0.1
Small	25	0.2
Medium	50	0.4
Large	100	0.8
Very Large	250	1.5
Huge	500	4.0

TABLE 3.1: Population sizes.

In the beginning ($t = t_0 = 0$) there are n_0 buyers and n_0 sellers in the market. This is to ensure that negotiations can start from the beginning. Over time, some buyers and sellers enter and some may exit. The numbers of entries for the parties are independent variables, but follow the same standard Poisson distribution, with the parameter $\lambda(t) = i\frac{t_{lastEntry}-t}{t_{lastEntry}}$, where i is the basic entry intensity, $t_{lastEntry}$ the last turn that entries are possible and t is the current turn. This formulation means that entries are more probable earlier in the experiment. This is realistic because the parties are more likely to find a contract if they enter early. This is especially true for the providers, because we assume that the provision of the service takes time and they cannot wait until the last moment to find a consumer. In the experiments we discuss in this thesis, we have $t_{delivery} = 1000$, $t_{lastEntry,s} = 800$ and $t_{lastEntry,b} = 900$ and we use different population sizes as described in table 3.1.

3.2 The Negotiation Parameters

We assume that the provisioning costs money (requirement **R1**). Specifically, in order to provide the service at the delivery time, the provider s has to invest a cost c_s at time $t_{c,s}(< t_{delivery})$. In the experiments we discuss in this thesis, $t_{c,s}$ is selected at random from either:

- *Any*: $Uniform(\max(0, t_{e,s} + 1), 1000)$,
- *Last Half*: $Uniform(\max(500, t_{e,s} + 1), 1000)$,
- *Turns 700-800*: $Uniform(\max(700, t_{e,s} + 1), 800)$, or
- *Last Negotiation*: $Uniform(\max(900, t_{e,s} + 1), 1000)$,

where $t_{e,s}$ is the time of entry for provider s. The time $t_{c,s}$ is selected independently for each provider using the same interval. Each provider has a quality q_s, which is selected at random from $Uniform(0,1)$. The cost is a function of quality and time:

$$C_s(q_s, t) = \begin{cases} 0, & \text{if } t < t_{c,s}, \\ c_s = 0.5q_s, & \text{if } t \geq t_{c,s}. \end{cases}$$

These provider characteristics are mapped into typical bilateral negotiation parameters by setting the reservation price r_s equal to the provider's preparation cost c_s and the deadline to $t_{c,s}$. This means that the provider will never accept a price that is less than its costs and that if the provider does not have a contract when it should start preparing for service, it will exit the market. The provider's utility for a contract is: $U_s(p, c_s) = p - c_s$, where p is the contract price.

The consumers do not have costs, but each consumer b has a deadline $t_{x,b}$, which is selected at random from $Uniform(t_{e,b}+1, t_{delivery})$, where $t_{e,b}$ is the time that b entered the market. The consumer's utility for the contract is $U_b(q,p) = V_b(q) - p$, where the value function $V_b(q)$ is:

$$V_b(q) = \begin{cases} 0, & \text{if } q < q_b^{\min}, \\ v(q), & \text{if } q_b^{\min} \leq q \leq q_b^{\max}, \\ v(q_b^{\max}), & \text{if } q > q_b^{\max}, \end{cases}$$

where $v(q)$ is the consumers' common value function and $q_b^{\min}$ and $q_b^{\max}$ are consumer specific parameters of that function. Here we assume simply that $v(q) = q$. This means that each consumer has a minimum useful quality $q_b^{\min}$ and any service that does not offer at least this is worthless ($V_b = 0$). On the other hand, the consumer also has a maximum useful quality, $q_b^{\max}$, which gives him his full utility. Any improvement above this level does not increase the value of the service to the consumer in question. The parameters $q_b^{\min}$ and $q_b^{\max}$ are selected for each consumer independently at random from $Uniform(0, 0.5)$ and $Uniform(0.5, 1.0)$ respectively. The consumer's reservation price for a given service is then equal to its value.

3.3 Decommitment

Since the original contract is always beneficial to both parties ($U_s > 0$ and $U_b > 0$) they would not consider abandoning it without some external force. We therefore introduce a possibility of an adverse impact that decreases the value of the contract to the party in question. This decrease may make the contract counter-productive to the affected party or parties and he or they may want to decommit. We use a_b and a_s to denote the probability that the buyer or seller (respectively) will be affected.

For the provider, the decrease means that the cost of providing the service increases by amount L_s and this will decrease its utility by the same amount. He will then need to make a decision on whether or not to decommit from the contract in this new situation. The decision is influenced by the decommitment fee f_s. We assume that the provider s will decommit at turn t if and only if:

$$\begin{aligned} U_s(\text{contract}|L_s = l) &< U_s(t_{decommit} = t) \\ p - c_s - l &< -f_s - C_s(q_s, t). \end{aligned}$$

where $U_s(t_{decommit} = t)$ is the seller's utility, when he decommits at turn t and l is the amount the utility decreases. Here we use the following ten values $l \in \{0.1, 0.2, ..., 1.0\}$. So, the seller decommits if the decreased utility is lower than the cost it has already paid and the decommitment fee it has to pay to get out of the contract. The seller learns of the loss at some point t_l (selected at random) between the time the contract was formed $t_{contract}$ and the time it was due to be performed ($t_{delivery}$) excluding both of the extremes. However, we assume that this loss itself is always avoidable, if the contract is abandoned before it is delivered. This means that the additional cost has to be paid just before the delivery. It is not possible that this additional cost is paid if there is no delivery.[2] There can only be one effect per party and the effect is always final. All possible moments for learning of the effect are equally likely, i.e. $t_l \sim Uniform(t_{contract} + 1, t_{delivery} - 1)$. If $t_{contract} + 1 > t_{delivery} - 1$ we assume that there can be no loss.

The same applies to the buyer, except for two important differences. For the buyer, the impact l decreases the value of the contract ($V_b(q)$). The fact that the effect can be avoided also here by decommitting at any time before the deadline is clear. However, because the buyer can, in many types of service, just ignore the service delivered, the value cannot be enormously negative. We therefore assume that the impacted value cannot be lower than -0.05 ($V_b(q) \geq -0.05$). This small negative value would then come from accepting the service and disposing of the results. This means that the utility of the buyer can never go below $-0.05 - p$. We

[2]However, it is possible that the seller has to pay the original cost even without a delivery as explained earlier. Only the extra cost is tied to the actual performance and will always be avoided, if there is no performance.

do not make a similar assumption with the seller, because the cost of producing the service can (in theory) increase without any limit (hardware failures, resource shortages and strikes can make the service very expensive to perform). The second difference is that for the buyer, always $C_b = 0$ (for the seller $C_s \geq 0$).

Since, in this chapter, we are interested in the system-level performance and common good, we use the sum of utilities of all buyers and sellers in the market as a performance measure. We chose the sum of utilities because it is the simplest way to measure common good. In addition, it is also the measure used in the law and economics literature.[3]

The total utility for a buyer and seller pair in the case the seller decommits from the contract at turn t is:

$$\begin{aligned}
&U_{b+s}(t_{s\ decommits} = t)\\
&= U_b(t_{s\ decommits} = t) + U_s(t_{s\ decommits} = t)\\
&= f_s + (-f_s - C_s(q_s, t)) = -C_s(q_s, t).
\end{aligned}$$

Similarly, the case where the buyer decommits $U_{b+s}(t_{b\ decommits} = t) = -C_s(q_s, t)$. In case both parties decide to decommit at the same turn, we assume that $f_s = f_b = 0$. It is clear that the total utility is equal to $C_s(q_s, t)$ here as well. The affected parties avoid the utility decrease of the contract, because there is no contract any more, but the decommitter will have to pay the fee (f_s or f_b) to the victim.[4]

If the parties decide to perform the contract despite the utility decreases, the total utility is:

$$\begin{aligned}
&U_{b+s}(\text{contract}|L_s = l_s \& L_b = l_b)\\
&= U_b(\text{contract}|L_b = l_b) + U_s(\text{contract}|L_s = l_s)\\
&= (V_b(q_s) - p - l_b) + (p - c_s - l_s) = V_b(q_s) - c_s - l_b - l_s.
\end{aligned}$$

A *decommitment policy* is a set of rules that specifies the amount the decommitter (the party decommitting) should pay to the victim (the decommitter's opponent) in case of decommitment. We will discuss several decommitment policies in this

[3]In our earlier work (Ponka and Jennings 2007), we used an expected utility of all contracts as a performance measure. However, in the more complicated situations we investigate in this thesis, calculating the expected utility would be much more difficult and we have instead selected the time for adverse impact (potential decommitment time) at random.

[4]Note that the fee does not change the total utility, just the distribution of wealth in the society. However, as explained earlier, the fee can affect *when* and *if* contracts are decommitted from and that can have an impact on the welfare of the society.

work, some of which are only to be used in a certain settings for a certain reason, but we will start by discussing shortly some basic policies that have been used in the literature (or can be easily derived from such policies). The common factor for all these policies is that they are not environment- or issue-specific. These policies are simple and usually they are used when an option of decommitment is needed, but the role of the decommitment policies and fees are not considered in detail. Specifically, we consider the following:

- *Not Allowed*: The contracts are absolutely binding and decommitment is not possible.
- *Constant*: The decommitment penalty f is constant; here we investigate cases where $f \in \{0.00, 0.25, ..., 1.00\}$.
- *Increasing*: The decommitment starts with min at t_{min} and increases linearly to max at time t_{max}. We investigate cases where $min = \{0.00, 0.25, 0.50\}$ and $max = \{0.25, 0.50, ..., 1.00, 1.50, 2.00, 2.50\}$ and $min < max$. There are three variations (all with $t_{max} = t_{delivery}$):
  - *Contract Time Only*: $t_{min} = t_0$ and $t = t_{contract}$.
  - *Decommitment Time Only*: $t_{min} = t_0$, and $t = t_{decommit}$.
  - *Both*: $t_{min} = t_{contract}$ and $t = t_{decommit}$.
- *Constant Price* (Andersson and Sandholm 2001): The decommitment fee is a fraction of the price (p). Here we investigate cases where $f = \{0.5p, 1.0p, ..., 2.5p\}$.
- *Increasing Price*: This has the same variations as the increasing policy (contract time only and decommitment time only variations were used in (Andersson and Sandholm 2001)), but the minimum and maximum are fractions of the contract price. We investigate cases where $min = \{0, 0.25p, 0.5p\}$ and $max = \{0.5p, 1.0p, ..., 2.5p\}$ and in all cases $min < max$.

In total, this means that there are 107 variations. Nevertheless, there is still a large number of possible policies and an infinite number of parameter values that we do not investigate. We have tried to choose a reasonable sized selection of the most obvious policies that have been used in literature or that are quite straightforward variations of those policies. All of these policies are non-compensatory, which means that they do not try to compensate for the losses of the opponent. We will also introduce two simple compensatory policies that explicitly try to compensate for the losses of the victim, namely:

- *Expectation Damages*: The fee is the opponent's expected profit (his utility if the contract is performed properly) plus his costs at decommitment time.
- *Reliance Damages*: The fee is equal to the opponent's costs at decommitment time.

We will also discuss some situation-specific variations of these policies. These basic policies rely on complete information about the opponent's losses but in some cases, we will also discuss some other policies that try to mimic the effect these policies have, but under incomplete information. We will now turn our attention to the specific settings and decisions.

3.4 Implementation

The agents and the marketplace were implemented in Java (Sun Microsystems' J2SE 5.0) using Eclipse 3.2.1 as a development environment.[5] No agent frameworks were used. The only external Java library used is Apache's log4j (version 1.2.13)[6], which allows management of log files. Our marketplace software uses a configuration (text) file that specifies the setting to be run in great detail and the software configures itself accordingly at run-time (before the experiments start). Around 200 values are calculated for each and every run and these values are printed into a text file along side with the configuration used to generate the results and a timestamp. The experiments were run mostly on the university's Beowulf cluster Iridis, [7] although some experiments were also run in a desktop and a laptop running Windows XP SP2.

After an experiment, the text file generated by the software was analysed in two stages. First, a simple Perl program (utilising an external Statistics::Descriptive module[8]) is used to calculate averages and standard deviations for each value. Then the actual statistical analysis is done using these values with Microsoft Excel.

[5]See `http://java.sun.com/javase/downloads/index_jdk5.jsp` for J2SE5.0 and `http://www.eclipse.org` for Eclipse.

[6]`http://logging.apache.org/log4j/1.2/download.html`.

[7]`http://www.southampton.ac.uk/isolutions/computing/hpc/iridis/`

[8]`http://search.cpan.org/dist/Statistics-Descriptive/`, version 2.6.

Chapter 4

The Performance Decision

In the *performance decision*, the party has to decide whether or not to perform according to a contract that it has entered earlier. The alternative is to decommit from the contract and pay the other party a decommitment fee and all parties are expected to take the alternative that maximises their personal utility. As discussed in the literature review (section 2.2.3.1), the optimal policy for the society is to set the decommitment fee equal to the losses of the victim. That way the decommitment occurs if and only if it benefits the society. However, the situation is not as simple as that. There is a question of how the opponent's loss is defined in different situations and a question of how this can be made to work under incomplete information, when the losses are not known exactly and the victim cannot be trusted to give an honest estimate of them. We will discuss these issues in this section. The work described in this chapter corresponds to our research contribution **C1**.

We will first discuss the changes to the basic model to better investigate performance decision and to try our ideas about it (section 4.1). This includes four different settings and the compensatory decommitment policies that we have devised for them (section 4.2). We will then discuss the empirical evaluation of our ideas and policies. Finally we will conclude with a summary of our findings (section 4.3).

4.1 The Problem and the Modified Models

In this work, we will focus on situations in which the parties always exit the market after they have found a contract. However, we investigate four different settings by varying two parameters: possibility of re-entry and the number of possibly affected parties. The first one, *possibility of re-entry*, determines whether, after the contract has been decommitted from, the victim is allowed to re-enter the market to find a substitute contract for the one it lost. This is typically possible in the case of simple and standardised services that can be provided or consumed by many other parties. In case of standardised services, it makes sense for the buyer to try to find another provider for his service or the seller to try to find another consumer, who would be interested in exactly the same service. In more customised service settings, re-entry is deemed to be a waste of time and other resources, because other consumers cannot readily use the service or other providers cannot provide it.

The second factor that we will use to vary our settings is *the number of potentially affected parties*. This means that the adverse effect mentioned can potentially influence either one or both of the parties. When only one of the parties can be affected (either $a_b = 0$ or $a_s = 0$), the possibility that the potential victim will be affected as well need not be considered. However, when it is possible that both parties are affected, the parties should take this into account. Nonetheless, we will assume that the parties will always decommit as soon as the need arises and therefore do not engage in any type of strategic behavior in this regard. Strategic behaviour could occur when there is a reasonable chance that the opponent would want to decommit, and it would mean that a party would wait for the opponent to decommit first to avoid paying the decommitment fee (and actually receiving one from the opponent). This problem was discussed in Sandhom and Lesser (2001), but will not be addressed in our work. We will, therefore, discuss four different settings in turn:

- one party potentially affected, no re-entry (section 4.1.1)
- one party potentially affected, re-entry possible (section 4.1.2)
- both parties potentially affected, no re-entry (section 4.1.3)and
- both parties potentially affected, re-entry possible (section 4.1.3)

4.1.1 Compensatory Policies in Markets with One-Sided Decommitments and No Re-Entry

As discussed in the literature review (section 2.2.3.1), an important aspect of compensatory policies is how the victim's loss is measured. In law, the compensation usually aims to put the victim in the same financial position as when the contract had been performed as agreed. In a basic case, this means compensating not only for the victim's costs, but also for his expected profit. A counterbalance of this quite extensive liability is the victim's duty to mitigate his loss. We will discuss compensatory policies in two parts: (i) those that have access to complete information about the opponent's profits and costs and (ii) those where this information is incomplete.

In the *complete information case*, we can use the decommitment policy that is optimal according to efficient breach theory:

- *Expectation Damages*: The fee is the opponent's expected profit (his utility if the contract is performed properly) plus his costs at decommitment time.

The compensation of costs is important. If the costs of the opponent are not compensated, the opponent would either have to take a risk that his costs could be wasted in case of decommitment or not to enter a contract at all. If he was willing to take a risk, he would sometimes incur losses. On the other hand, if he was to stop providing a service that is useful to the society, the society would suffer. Compensation of costs can therefore be seen as essential to contracting in situations where there is a delay between agreement and performance. In many situations the cost can be a significant portion of the contract price.

However, the costs should only be compensated when they could not have been avoided. The relevant legal principle here is *duty to mitigate*. Once the breach (or decommitment) is a fact, the victim is expected to take reasonable action to mitigate his losses. This duty is enforced by allowing damages only for losses that he could not have avoided this way (Treitel 2003). In this simplest setting, the duty to mitigate means that if there is no reason for the seller to finish the preparation for the service that the buyer does not want, the seller should not do so but if he does, he should not be compensated for it.[1] The provider should therefore exit the market before his costs are due.

On the other hand, compensating for the costs is not enough, because each contract also means profit for its parties and these profits increase the total utility. If the expected profit is not compensated for, it is possible that decommitment occurs when the increase of the decommitter's utility is lower than the expected profit that the victim loses. To investigate this, we use a decommitment policy that compensates only for the costs. This is typical in tort law and is called reliance damages:

- *Reliance Damages*: The fee is equal to the opponent's costs at decommitment time.[2]

In the *incomplete information case*, the situation is more complicated. There are many possible ways to create a compensatory decommitment policy that works under incomplete information, but that aims to compensate for the loss the decommitment causes. Here, we introduce the most obvious one: the analytic compensatory. In this policy, the victim's loss is estimated analytically using available information. In other words, the policy uses estimates of the cost and value functions, the distributions for the deadlines, and so on, to analytically estimate the loss. The accuracy of these estimates can vary, but here we will only consider the most accurate setting. In our earlier work (Ponka and Jennings 2007) we showed how the quality of information affects the performance of this policy.[3]

- *Analytic Compensatory*: The fee is equal to the expected loss for the victims in similar circumstances.

In more detail, the decommitment fee for the buyer is:

$$f_b(p, q, t) = D(t)p + (1 - D(t))(p - EC(q)),$$

[1]The actual legal situation is usually a bit trickier and requires that the seller has accepted that the buyer doesn't want the service. This is handled under the doctrine of anticipatory breach and will be ignored here. Since the parties are in the market where leveled decommitment contracts are used, we assume that they always accept decommitments. Such acceptance might, for instance, therefore be a requirement for entry to the market.

[2]This policy was used in Excelente-Toledo et al. (2001), but it was called sunk costs.

[3]In that work, we also used another compensatory policy, the *Expected Loss*, which used empirical data instead of analytical estimates. We showed there that the better (more accurate) information usually improves the performance (the total utility of the market players).

where q is the quality, $EC(q)$ the estimated cost function and $D(t)$ a probability that the seller has paid the cost at turn t. In a similar fashion, the fee for the seller is:

$$\begin{aligned} f_s(p,q) &= EV(q) - p \\ &= [F_{mid}(q) \cdot V(q) + F_{max}(q) \cdot q_{max}(q)] - p \end{aligned}$$

where $EV(q)$ is the estimated value function, $F_{mid}(q)$ is the probability that the quality q is between the buyer's minimum and maximum value, $F_{max}(q)$ is the probability that quality q is above the maximum quality, and $q_{max}(q)$ is the estimated value for $q_b^{\max}$ for the opponent b in the latter case. In case this is negative, the fee is zero.

In the other compensatory policy, average loss, we use use the notion of 'normal' or typical loss. In law, the unusually large losses (even if real) are not compensated if they were not foreseeable to the other party.[4] In law, this limits the maximum for damages, but here we use it as the measure of loss. We compensate the typical loss for the opponent in the same circumstances. The circumstances are determined by information that is available to both parties, such as the turn the agreement was reached ($t_{contract}$), the decommitment turn ($t_{decommit}$), the quality of the service (q_s) and the contract price of the decommitted contract (p). As the measure of typical loss we use the average:

- *Average Loss*: The fee is equal to the average loss for the victims in similar circumstances.

Here, we investigate the variation in which the similarity of the situation is assessed by the contract price, quality, contract turn and decommitment turn, and in which each of the first three factors are divided into k categories and we have accurate information of all possible decommitment turns; that is, we have $1001k^3$ different categories. For our experiments, we establish the typical loss simply by running the market 1000 times in advance and by calculating the average losses experienced by the parties in different situations (from each contract we get the information on losses of all possible decommitment situations). In the calculation of the fee, the situation is first categorised in terms of all three factors and the similar situations are those that belong to the same categories in all three factors.

We assume that the decommitment fees are established by the marketplace and the fee for any set of circumstances is always known by all parties. This policy

[4] The actual rule is that the loss must either be directly or naturally following from the breach or the other party knew or should have known of this loss at time of signing of the contract (Treitel 2003). This rule can be also seen as a special case (market-set) of liquidated damages.

gives the victim incentives to minimise his loss, because the compensation is set in advance and hence all the savings the victim can make are going to benefit him (and the society).

4.1.2 Compensatory Policies in Markets with One-Sided Decommitments and Re-Entries

When the victim can re-enter the market and try to get a new contract to replace the lost one, it is clear that there are two possible outcomes. On one hand, the victim may be able to find a reasonable substitute, or on the other hand, he might fail to do so. The latter is relatively easy as it is the same as the case with no re-entry. However, it may be much more difficult to estimate the probability of finding a new contract and what such a contract will be.

In law, the duty to mitigate damage means that if it is possible and reasonable, the victim is expected to make a substitute transaction in the market. In other words, a buyer that does not get the product from one source is expected to acquire it from another and a seller that is unable to sell the product to one buyer is often expected to try to sell the product to another customer. If the victim is successful, his loss is then the price difference between the new and original contract.[5]

Accordingly, the prima facie rule on contract law damages, *market price rule*, says that, if there is a market in which a victim can make a substitute transaction, the damages are limited to the difference between the contract price and the market price (Atiyah et al. 2005). In theory, the actual contract price for the substitute transaction is irrelevant and only the market price matters. However, in many settings it may be difficult or even impossible to determine one market price in a given market at a given time. In such situations, if the victim can show that he has acted reasonably, the market price will probably be held to be equal to the actual contract price. Unlike in the previous setting, here it is possible that in case of a substitute contract, the loss of the victim may actually be negative (i.e., he may actually benefit from decommitment). This occurs if he is able to find a better contract than the one that he had. In law, this situation is managed by awarding no damages. Accordingly, we require the decommitment fee to be non-negative and in case the victim benefits from the decommitment, the fee will be zero.

[5]Stricly speaking also reasonable costs that are related to securing a substitute contract should be compensated. Here we assume such costs are neglible (= 0).

In our setting there is no one market price at any given time, since different negotiations on the same service can end in a different result, even at the same instant. Therefore, in our market, the actual price of a substitute contract should be the starting point for loss assessment. However, the market price rule assumes that the substitute purchase is of similar (preferably identical) quality and therefore the loss is directly the price difference. Since in our setting, the quality of the providers and value function of the consumers vary and the matching occurs at random, this is not directly applicable to our setting. We therefore modify this principle to also take into account the differences in quality (value). The loss is then the difference between the utilities of the new and the old contracts.[6]

What makes this variation very different from the case without re-entry is that outside certain clear cases (where there are many consumers/providers with very similar tastes/prices and therefore the probability of a replacement contract is close to one and the price is known), the outcome is genuinely uncertain. Nobody, not the decommitter or even the victim, will know for certain what will happen before the decommitment occurs (ex ante) and the victim is forced to find out. On the other hand, after the dust settles (ex post), it is possible to assess what the loss actually is and therefore what the decommitment fee should have been. This uncertainty holds even for the market rule in law. If the market is not fully developed and the availability of substitute contracts and the possible prices/qualities is not known, nobody knew what was going to happen. The non-performance (decommitment) occurred because one of the parties (the decommitter) estimated that he is better off not performing and paying up whatever the court finds him to be liable for. This means that even the legal rules for damages can lead to inefficient non-performances or performances under incomplete information: the decommitter's estimate could be wrong. It might underestimate the loss, in which case he might not perform even when it would be in his (and the society's) interest, or it might overestimate the loss, in which case he might perform even if it would have been better for him and for the society to abandon the contract.

The problem is that in cases of genuine uncertainty, the inefficient (adverse) non-performances and/or performances cannot be completely avoided. The decommitter has to make his decision based on the information he has and even if he had very good estimates for the two cases (success or failure in finding the substitute contract) and even for the probability of success, but in any given case, he still does not know which one is going to happen. Now, the legal rule puts the risk of making

[6]Unless the new contract provides better utility and, as explained, the loss would be considered zero.

the wrong decision on the decommitter. The victim will always be fully compensated for whatever loss he incurs. The fee (actual loss) can therefore be higher or lower than the estimated loss (actual loss model). Another approach would be to use average losses or some other fixed values as damages or decommitment fees. In this case, the decommitter would always know what his final liability will be (estimated loss), but the victim does not know his actual loss until he is forced to find out (estimated loss model). Here, the victim can be overcompensated but also undercompensated, so he takes some of the risk.

The common good is not directly affected by the choice between these two alternatives, since in both cases the potential decommitter uses the estimate and after the decision to decommit has been made, the difference only affects the distribution of utility between the parties. However, as explained, the choice does affect the distribution of risk between the parties. On the other hand, the actual loss model suffers from the problems with incentives. Namely, it may be difficult to ensure that the victim who has been unable to find a good substitute contract has failed despite his best efforts and not because he preferred to take the full expectation damages instead. This difference can be limited somewhat if the actions taken by the victim can be reviewed[7], but a potential problem still remains. The expected loss model does not have this problem because finding the best possible substitute contract is in the victim's interest. However, the problem here is that sometimes even the best substitute available with the decommitment fee may not be enough to secure the victim the same utility as the performance of the original contract would have produced.

As discussed in the introduction, in part I, we will be interested in total utility, not the individual parties' utilities. We use the estimated loss approach, so that our agents always try to find the best possible replacement contract after decommitment. However, given the discussion above, our results would not be that different if the actual loss approach was used (assuming the incentivisation problem could be avoided). On the other hand, we do not expect the victim to try to find as good a contract as the one that was decommitted. Instead, we assume that the reasonable course of action is to re-enter the market and try to find a contract as if there was no earlier contract at all and that all players will always do so. This is

[7]This means that an unhappy decommitter could get some trusted third party (market arbitrator) to review the actions that the victim took after the decommitment and assess if they were reasonable in the circumstances. If the arbitrator believes that the victim did not negotiate properly or turned down perfectly reasonable offers, it could decrease the fee accordingly. This is similar to court review of the damages, but could probably be done in relatively simple manner, even automatically.

a reasonable assumption, because in our market, the victim has a limited number of chances for getting a substitute contract and the contract may not be possible at all in some matches, so the prudent (risk-averse) course of action is usually to take whatever is on offer. Thus, any contract that provides positive utility for the victim is better than no contract.

As in the previous case, we consider first the complete information case and then the incomplete information one. *Complete information* in this setting means that the outcome of each and every decommitment of any contract at any given moment is known at the potential decommitment time. In other words, the parties know if the victim is going to find a new contract and what that contract (if any) is going to be like:

- *Optimal Market Rule*: If the victim finds a substitute contract, the compensation will be the utility difference between the two contracts (or zero, if the substitute contract is at least as good as the decommitted contract). If the victim does not find a substitute contract, full expectation damages (expected profit + costs) are rewarded.

The *Optimal Market Rule* policy will therefore always reward the actual losses (ex post). As just explained, in most realistic settings (in which the outcome for the victim is not known in advance), we have to settle for estimating the losses ex ante.[8] As just discussed, even the contract law does not assume such information and a more indirect reasonableness test for the victim's actions is used. However, there are situations in which we can get very close to the Optimal Market Rule. This can occur, for example, when there is an efficient market (when there are many other providers/consumers and quite similar prices). In such cases, the outcome of decommitment is known to a very high degree of accuracy. This can occur in some real-life markets, although it must be considered more of an exception than a rule.

In practise, to simulate the *Optimal Market Rule* we run the market and take note of all the losses, run the market again setting the decommitment fees equal to those losses for the turns that decommitment occurred in the previous run. We repeat this procedure until the fees match the actual losses exactly. We also use some simple heuristics to minimise the number of iterations we have to do. For

[8] And as also explained, once such estimates are used, in theory, it matters not if the fee is set according to the actual or estimated losses, because the decommitment decision is always based on estimates and the rest is just distribution of welfare.

example in many cases it matters little if the decommitment occurs in turn 192 or 193 if the parties will be affected at turn 349 and 493 respectively and the seller's cost will be paid in turn 483.

Creating a compensatory decommitment policy for the markets where re-entries are possible and *information is incomplete*, is more challenging than for the case without re-entry. As explained earlier, this is because the victim's re-entry can either end in him finding a replacement contract or him failing to do so (see figure 4.1). In addition, the case where no substitute is found is divided into two subcases according to whether or not the opponent has accrued costs before the decommitment occurs. In order to have effective estimates of the overall loss, the following information is needed:

- the probability that the victim will find a substitute contract (branch a),
- the probability that the victim has paid its cost (branch b),
- the expected utility in case the victim finds a substitute contract and
- the expected utility in case the victim does not find a substitute contract

If the cost has not been payed, the victim should exit the market before he has to pay the cost to mitigate his loss. On the other hand, if he has paid the cost already when the decommitment occurs, the victim should stay in the market until the time of performance trying to find a new contract. For this case, we use:

- *Semi-Analytic Market Rule*: The fee is equal to the difference between the estimated current utility and the expected utility given the decommitment. The latter is calculated by empirically estimating the probability that a replacement contract can be found (given the decommitment time and quality of the contract, branch a) and the utility of any such contract (again given quality). The utility for the failure is calculated analytically as in the *Analytic Compensatory policy.*

The *Semi-Analytic Market Rule* uses information about the current contract to estimate the opponent's loss and considers the structure of the problem (successes and failures) to assess the expected loss. To determine the success probability and the expected utility of success, it relies on empirical results that were gathered by running the market with exactly the same parameters for 1000 times and calculating averages.

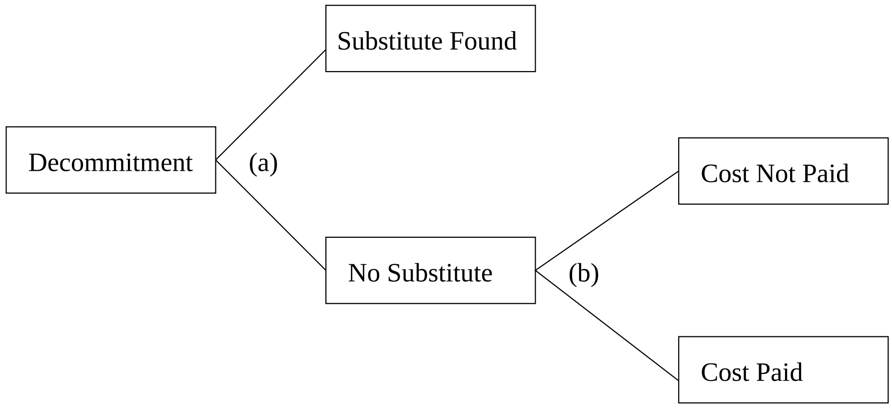

FIGURE 4.1: Possible outcomes of decommitment for the victim.

4.1.3 Compensatory Policies in Markets with Potentially Two-Sided Decommitments

In addition to the possibility of re-entry, the other factor that we vary here is whose utility can be adversely affected after the contract has been formed. Here we discuss the case where the effect can decrease the utility of both parties. This means that it is now possible that the potential victim has also been affected. This will change the situation in two ways. First, it is possible that both parties will want to abandon the contract as soon as the effect occurs. We will therefore adjust all our policies so that in case both parties want to decommit at the same time, the fees are cancelled and therefore nobody pays anybody anything. Second, this will mean that the possible utility loss of the opponent might have to be considered when assessing his loss.[9]

So, the possibility of the opponents' utility losses should be taken into account in the decommitment fees. To see why this is so, consider a simple example. Let $V_b(contract) = 0.75, p = 0.50$ and $c_s = 0.25$. Now, let both parties be affected by a utility loss of 0.3 and using decommitment policy *Expectation Damages*. This means that:

$$
\begin{aligned}
U_s(stay|l_s = 0.3) &= 0.50 - 0.25 - 0.30 = -0.05 \\
U_s(decommit) &= -U_b(stay) = -0.25 \\
U_b(stay|l_b = 0.3) &= 0.75 - 0.50 - 0.30 = -0.05 \\
U_b(decommit) &= -U_s(stay) = -0.25
\end{aligned}
$$

[9]As mentioned earlier, the chance that the opponent might want to decommit as well, might lead to strategic considerations or waiting for the opponent to decommit first. We do not consider this problem in this work, although by taking the opponent's utility loss into account, we think this problem can be decreased (but not completely removed).

In both cases, the utility of staying in the contract is higher than decommitting, although both parties and the society are worse off because of the contract. The key would be to take the utility loss into account. That will decrease the decommitment fee to zero[10] and both parties would decommit, which was the efficient result here.

Also this adjustment and principle behind it can be found in contract law. Because the aim of damages is to compensate for the loss and to put the party into the financial position he would have ended up in if the contract had been performed as agreed, his overall position is taken into account (Treitel 2003). The relevant factors include any benefits which he may have obtained under the broken contract, and his release from obligations under it. In this setting, it would mean that the seller is released from having to pay whatever costs he has not paid yet. This will always include the seller's last-minute costs (adverse effect) if applicable. On the other hand, the fact that the buyer's valuation of the service has decreased should also be taken into account, because it diminishes his profit. The courts will not generally order the defendant to pay an amount which will actually make the claimant's position better than it would have been if the contract had been performed.

In the *complete information case*, we simply add the opponent's effect to *Expectation Damages* and *Optimal Market Rule* policies:

- *Expectation Damages (with Opponent Effect)*: The fee is the opponent's expected profit (his utility if the contract is performed properly) plus his costs minus any adverse effect at decommitment time.
- *Optimal Market Rule (with Opponent Effect)*: If the victim finds a substitute contract, the compensation will be the utility difference between the two contracts (or zero, if the substitute contract is at least as good as the decommitted contract). If the victim does not find a substitute contract, full expectation damages (expected profit + costs − effect) are rewarded.

In both cases, the adverse effect is simply subtracted from the decommitment fee.[11] This is because the adverse effect always decreases the victim's profit and because the effect can always be avoided in full by decommitting a contract any time before the delivery.

[10] $U_b(stay|l_b = 0.3) = U_s(stay|l_s = 0.3) = -0.05$ and since the fee cannot be negative, it would be zero.

[11] If this adjustment would make the fee negative, the fee is zero.

A very similar adjustment is required to all the policies *under incomplete information* that were discussed earlier. We investigate three cases:

- *Opponent Effect Not Considered*: The possible decrease in the opponent's profits are not taken into account at all,
- *Opponent Effect Partially Known*: The amount of potential utility loss and its probability are known, but it is unclear *if* the loss has been incured by a specific opponent and, if so, *when* it becomes known to the opponent.
- *Opponent Effect Fully Known*: The actual amount of utility loss for a given opponent is known exactly.

The first case is the same as the *Semi-Analytic* policy that was discussed in section 4.1.2.

The effect occurs to the seller with probability a_s and to the buyer with probability a_b. In the partial information case, we assume that these probabilities are known, but the fact if a certain participant has been affected is only known to himself. All points of time between the contract time ($t_{contract}$) and delivery time ($t_{delivery}$) are equally likely, so the probability that the effect has become known at turn t is:

$$P(\text{effect occurred at turn } t) = a_s \frac{t - t_{contract}}{t_{delivery} - t_{contract}}$$

This is then multiplied by the effect to get the estimated effect. This estimated effect is then subtracted from the fee from those given by the *Semi-Analytic Market Rule* policies. To separate this policy from the other variations, we call it a *Semi-Analytic Market Rule (with Estimated Opponent Effect)*. The case where the opponent effect is fully known (as in the full information case) is called *Semi-Analytic/Simple Market Rule (with Opponent Effect)*.

4.2 Empirical Evaluation

Having introduced the various decommitment policies, we now compare their performance in the various settings of our base model. This section consists of three parts. First, we discuss our hypotheses (section 4.2.1). Second, we explain our experimental set-up and how the analysis was conducted (section 4.2.2). Finally, we discuss the actual results (section 4.2.3).

4.2.1 Hypotheses

The first hypothesis considers the circumstances in which allowing the victims to re-enter the market after decommitment can be useful. If there is no possibility of re-entry, the parties are always left without a contract in case of decommitment. This means no benefit to the society (the total utility of parties is $-C_s(t)$, the seller's cost at decommitment time). In contrast, if a re-entry is allowed and the victim has many opportunities to find a replacement contract, such a contract is likely and will contribute towards the welfare of the society. However, if the probability of decommitment in the replacement contract is also very high, the re-entry might be useless or even harmful. This is because instead of one loss $-c(t)$, there may be several of them. On the other hand, when the probability of decommitment is very low, the differences are less likely to be statistically significant. Given all this, we contend:

> **Hypothesis 1.** In cases where there is a good chance of finding a replacement contract and decommitments occur frequently, but not all the time, allowing re-entry for the victim improves total utility.

Our second hypothesis considers the potential utility loss of the opponent in decommitment fees. We expect that taking this loss into account is useful, because the utility loss of the opponent will decrease the utility of the contract to the society and the fee should reflect that. Thus we contend:

> **Hypothesis 2.** When both parties can decommit, taking the possible utility loss of the opponent into account improves the total utility.

The next two hypotheses deal with the performance of compensatory policies compared to their non-compensatory counterparts. In all four settings we discuss in this thesis, we first introduced the compensatory policies under complete information. In doing so, we explained that the basic idea of these policies is to put the decommitment fee to a level in which the parties decommit only when it is in the society's best interest and that this occurs when fees are equal to the victim's losses. The optimal policies use complete information to do this and therefore they should be better than any other policy in most circumstances. Of course when the utility loss is small, the best policy is often to stay in the contract and many non-compensatory policies (with high decommitment fees) achieve that as well, so the difference between optimal and good is small. On the other hand, when the utility loss is very large, it is often useful to abandon (almost) all contracts and some non-compensatory policies (those with low decommitment fees) will achive this as well.[12] Therefore, the optimal policies are likely to work best when the utility losses are intermediate and some, but not all, contracts become counterproductive.

However, there can be settings in which some usually non-compensatory policy can achieve a very good performance and, in some cases, the optimal policy and such non-compensatory policies yield a similar performance (no statistically significant difference). This is often because the non-compensatory policy happens to produce fees that are close to those of the optimal policy. The more complicated the setting, the less likely such a coincidence is, but when it occurs it can be very useful since specific non-compensatory policies are often much simpler than general compensatory ones. One extreme example of this would be *Constant Price (100 %)* policy for the buyer when the seller has no costs. Although it does not, in general, try to compensate for the loss, it is the same as the optimal strategy in cases where there is no re-entry or a possibility of two-sided decommitment. However, it will be sub-optimal if there are costs or if re-entries are allowed. Even in these cases, though, the optimal policy will not be worse than any non-compensatory policy. We therefore contend:

> **Hypothesis 3.** The optimal complete-information policy for each setting will yield at least as good a total utility as any of the non-compensatory alternatives and there are situations in which it will be better with intermediate utility losses.

Now, in most settings, complete information about the parties' losses may be difficult to obtain. We therefore have to consider policies that operate under incomplete information. Since the total utility is maximised when the losses are always perfectly compensated, we assert that in situations in which the losses can be more accurately estimated, and therefore the decommitment fees can be set closer to the optimal ones, we expect to see higher total utilities. To see why this is the case, we need to consider two ways in which the effect the policy has on individuals can differ from the optimal policy. First, the policy may force the party to stay in a contract even if the socially optimal action would be to abandon it

[12] However, typically the decommitment policies that achieve these two different goals are not the same ones. The policies with very high fees, which force parties to stay in contracts, are very usefully when the losses are small, but they perform very badly when the losses are large. And the policies that allow decommitment when the losses are medium or high may also allow decommitment when they are low and it would be useful to stay in the contract.

(adverse commitment). Second, the policy may allow the party to decommit from a contract even though it is still socially valuable (adverse decommitment). In the first case, the policy overestimates the loss and in the latter case it underestimates the loss. Both cases decrease total utility. Now, the closer the estimates are to the actual values, the fewer of these mistakes occur. The fewer mistakes, the less total utility is decreased and, hence, the higher it is. So, when the compensatory policies have access to relatively accurate information, the fees they set are likely to be closer to the optimal decommitment fees and therefore the society is likely to do better. Given this, we contend:

> **Hypothesis 4.** The best incomplete information policy for each setting will yield at least as good a total utility as any of the non-compensatory alternatives and with some intermediate utility losses it will be better.

Now, the whole point of using the compensatory policies is that they can improve the total utility of the market participants. The logic is that when these policies manage to more accurately compensate for the actual losses, the performance should be closer to the full information case. However, they do need sufficient information about the opponent's losses to achieve this. Speaking more generally, we believe that there is a clear relation between the accuracy of loss estimates and the performance (in terms of total utility). As a measure of performance, we use the average total utility over different utility loss cases. And as a measure of average compensation error, we take the average distance between the decommitment fee and the actual loss. For the full information optimal policies, such as Expectation Damages (in case with no re-entry) and Optimal Market Rule (in case with re-entry), this measure is zero, because the fees are always exactly actual losses. For other policies, this is a positive number. A similar calculation is performed to obtain the error for the buyer. Because of the method used to calculate the average compensation error, we cannot consider policies that do not allow decommitments at all. And because we use average utility over all possible utility loss cases, we cannot consider any policy that does not allow decommitments in any of the utility loss cases. Although this may sound something of a limitation, many policies do allow decommitments at all loss levels.

> **Hypothesis 5.** The smaller the average compensation error, the better the average total utility.

4.2.2 Experimental Setup

We ran the market with 10 different loss levels by setting l_s and/or l_b (depending on the setting) to each value of the set $\{0.1, 0.2, ..., 1.0\}$ in turn for each case and policy we investigated. In each case, we therefore had different contracts. We used 0.5 as the probability of adverse effect for the affected party ($a_b = 0.5$ and/or $a_s = 0.5$ depending on the setting) and we used the *Last Half* deadline setting for the providers unless otherwise stated. The population size for both parties is *Medium* (as explained in section 3.2) for the cases where both parties are affected. In cases where only one of the parties is affected, the affected population is *Large* in order to ensure that there are many opportunities for the victims if they are allowed to return to the market. Otherwise it might be that the victim returning to the market would be unable to find a negotiation partner and the return would often be unsuccessful.

We ran the market with the same setup 100 times and calculated the average total utility (and its variance) for all cases we investigated (different policies and different loss situations). We then conducted a simple two-sample t test to see if the two averages from the runs with different parameters were statistically different. Since we usually expect the compensatory policy to outperform the non-compensatory one, we performed one-tailed tests at each data point. We made the test between the compensatory policy in question and all non-compensatory policies separately. When we say that a compensatory policy outperforms non-compensatory policies, it means that it beats each and every one of them. In our graphs, we use the *Best of Non-Compensatory* line to indicate the best result that any of the 107 non-compensatory policies achieved in that data point. This line usually consists of several different policies.

Here and in all the experiments later in this work, we either compare two separate policies (tactics, strategies) or we compare one policy (tactic, strategy) to many others. In the latter case, this means that we compare the performance of this one policy to all the other policies in turn and we claim that the one policy is better than the rest if and only if it is able to beat *all* the other policies in these pairwise comparisons. This means that our one policy has multiple possibilities to fail to show a difference to the other policies and, therefore, we may *under*estimate the significance of differences in our experiments. In other words, we run somewhat increased risk of *Type II* error (accepting a null hypothesis when it should have been rejected). Assuming we had no differences between one policy and ten competing ones, the probability that our statistical tests would show a significant difference in a single test could be for example 5%. Now, in ten such tests, the probability that all ten tests show such a difference is only around $9.8 \cdot 10^{-14}$ and we are therefore less likely to make *Type I* error (rejecting a null hypothesis when it should have been accepted) but somewhat more prone to make a *Type II* error. For example, say the one policy is better than others, but with the probability of 5% our test fails to show it. The probability that we fail to show the difference at least in one of our tests is around 40% ($= 1 - 0.95^{10}$).[13]

However, the risk of the *Type II* error is decreased by the fact that often only a handful of competing policies is anywhere close to the one policy we compare them to and, therefore, the risk of a *Type II* error is very remote in most comparisons. It may still exist in some situations, though, and what all this means is that because we do not explicitly take this into account in our analysis, we are likely to *under*estimate the significance of differences in some cases and there might be differences sometimes when our tests show there are none. However, we usually get quite clear results and are able to draw our conclusions, so more complicated statistical analysis was deemed unnecessary in our work.

4.2.3 Results

We will now discuss the results in four parts. First, we will discuss the effect of allowing re-entry. Second, we will discuss the effect of taking into account the fact that also the opponent may have been adversely affected. Third, we investigate how our compensatory policies (under both complete and incomplete information) perform against the non-compensatory policies. And finally, we see if the average compensation error has any effect on the average total utility.

We start by investigating the effect of allowing victims to re-enter the market and to try to minimise their loss by finding a substitute contract. To this end, figure 4.2 shows the case where the adverse effect can decrease the utility of one or both parties at probability of 0.4. We can clearly see the difference between the two settings. The same policies do significantly better when re-entry is allowed and they settle clearly on different levels of performance when the loss is high. For both the *Expectation Damages* policy and *Best of Non-Compensatory* policies,

[13] This situation is then very different from the situation where we are interested in whether or not there is any differences between n policies and we have to do $n(n-1)$ tests to investigate. In this alternative setting, we do more tests and any positive result will mean a difference, which is why additional statistical methods are used in such situations to ensure the probability of *Type I* error is kept in check.

the difference is statistically significant with utility loss levels 0.2 and above (at $p < 0.001$ level). The same effect can be seen in all policies that allow a significant number of decommitments. Similar results can also be achieved in cases where only one of the parties is affected.

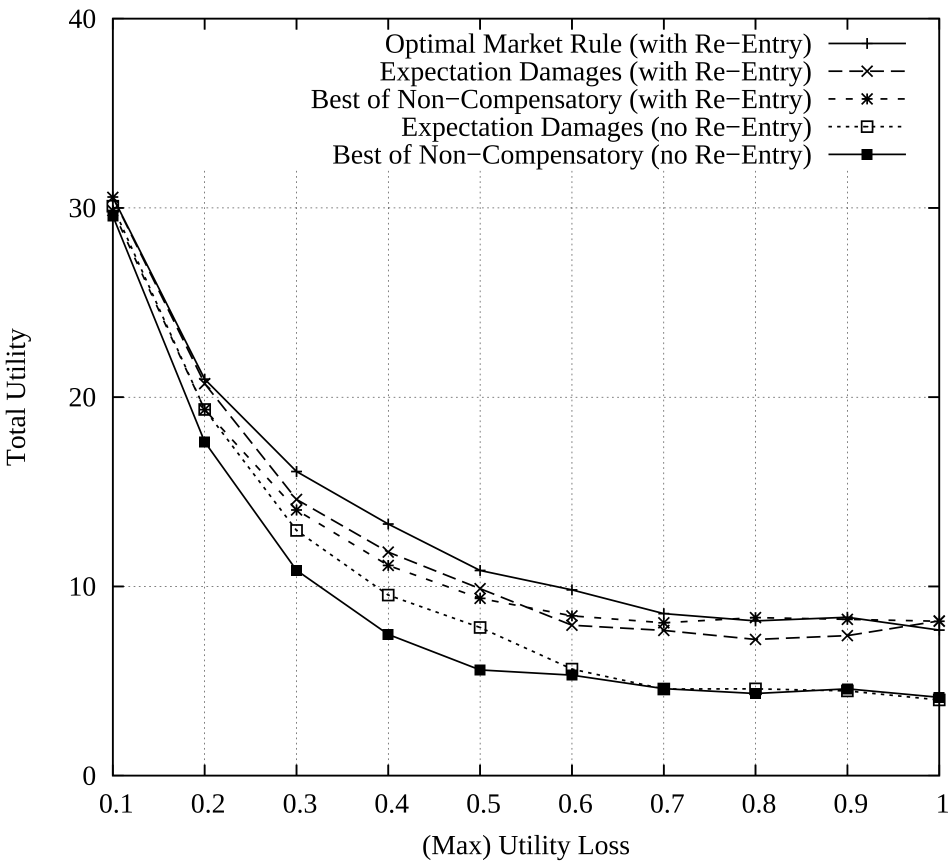

FIGURE 4.2: Effect of Re-Entry (Hypothesis 1).

As explained, this is because the victims can find a new contract that contributes towards the common good. In some cases, they can even do this several times before they find an opponent that will eventually perform it. So the re-entry is clearly useful. However, in hypothesis 1 we thought that re-entry is only useful when the probability of utility loss is not too large. To test this part of the hypothesis, we varied the utility loss probability $0.1, ...1.0$ and kept the loss itself at 0.5. The result of this experiment for the *Expectation Damages* policy with and without re-entry is shown in figure 4.3. As can be seen, the policy with re-entry dominates the policy without re-entry with low effect probabilities, but when the probability of utility loss increases above a certain level (0.6 in this case), the situation changes and the case where re-entries are not allowed starts to outperform the case where re-entry is allowed. Because we did not know where the switch occurs we did two-tailed t-tests for all data points. The policy with re-entry is better between loss levels $0.1 - 0.4$ (at $p < 0.001$ level) and at 0.5 (at $p < 0.01$ level). The change occurs at 0.6 where the policy without re-entry is better (at $p < 0.05$ level) and levels $0.7 - 1.0$ (at $p < 0.001$ level).

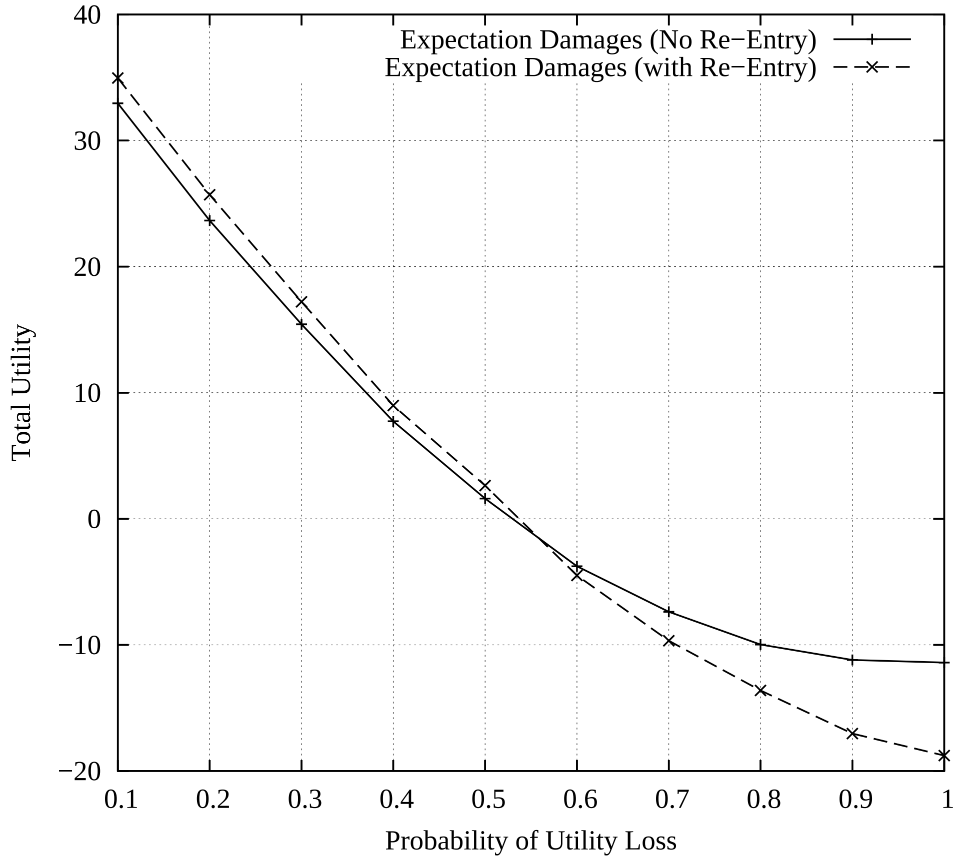

FIGURE 4.3: Effect of Utility Loss Probability on the Usefulness of Re-Entry (Hypothesis 1).

This is because even the successful re-entries (ones that find new contracts) are likely just to end in a new decommitment and in all decommitments, there is a chance that decommitment occurs *after* the provider has paid its cost and therefore re-entry actually ends up doing more harm than good. So, allowing re-entries is useful if the probability of utility loss is not too high, but can be harmful if the probability is very high. This is consistent with hypothesis 1 and we can therefore accept it.

We now turn to situations in which both parties can encounter an adverse effect that decreases their utility. Specifically, we try to determine if we should take the possible adverse effect of the opponent into consideration when we are setting the decommitment fees. We do this by running experiments in a setting with no re-entries and by trying two different complete information policies: *Expectation Damages* and *Expectation Damages (with Opponent Effect)*. The former does not consider the possible adverse effects of the opponent, whereas the latter does. Otherwise the policies are identical. From figure 4.4, the difference between the two is clear. There is no statistically significant difference with very low and very high loss levels. This is because in these cases, almost all the contracts are abandoned (high losses) or almost none of them are (low losses). But in the intermediate cases (the loss levels of $0.2 - 0.7$), the difference is statistically significant (at $p < 0.001$ level). This is because taking into account the opponent's decreased utility will make the decommitment fees compensate for the actual loss of the opponent at the time of decommitment and this allows parties to abandon contracts that are no longer useful for the society. In contrast, when the utility loss of the opponent is not considered, the fees are often set too high and the parties stay in contracts that have become detrimental to the common good. This of course only occurs in cases where both parties are actually affected. When there is only one-sided effect, the policies are exactly the same. Taking the utility losses of the victim is therefore useful.

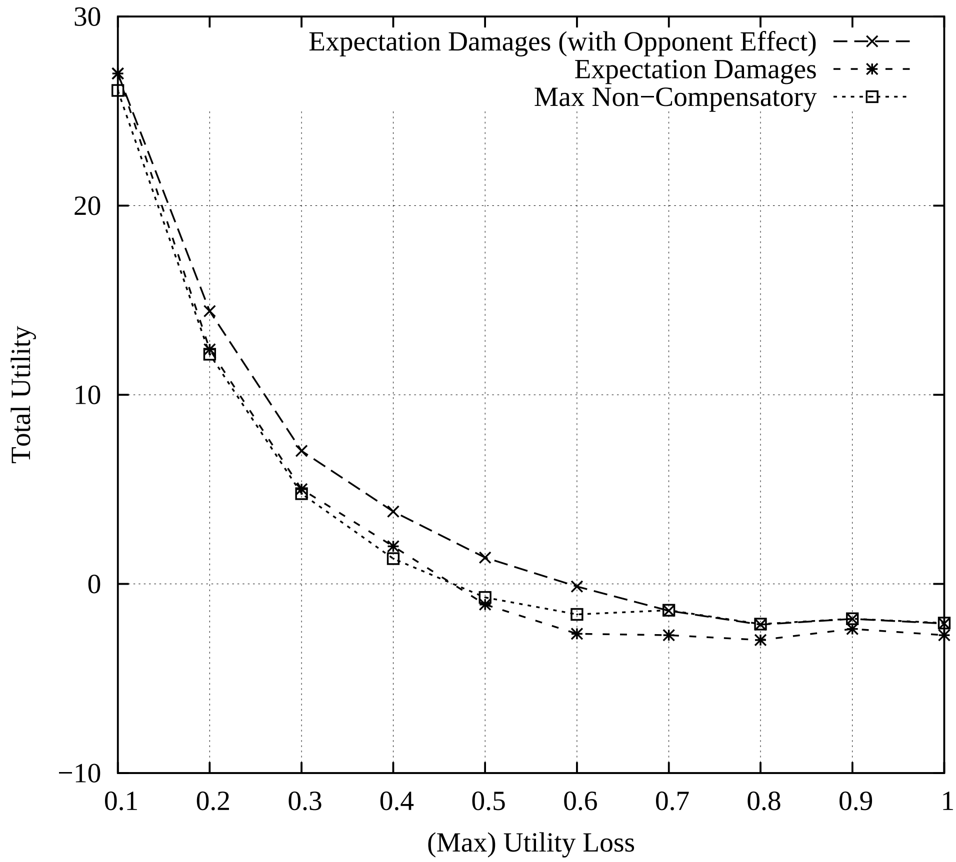

Figure 4.4: Effect of Opponent's Utility Losses (Hypothesis 2).

A similar effect can also be seen in cases where re-entries are allowed when there is incomplete information. In figure 4.5, we have three variations of the *Semi-Analytic Market Rule* policy with different stances and knowledge on opponent effects. We can see that the *Semi-Analytic Market Rule (with Opponent Effect)* (the one with the best information) outperforms the other Semi-Analytic Market Rule policies with utility loss levels $0.5 - 0.6$ (at level $p < 0.001$). Also the *Semi-Analytic Market Rule (with Estimated Opponent Effect)* policy is able to outperform the pure *Semi-Analytic Market Rule* (that does not take opponent effect into account at all) with loss level of 0.4 ($p < 0.01$ level). This shows that taking into account the utility loss of the victim is useful and that the better is the information about the loss, the better the performance in terms of total utility. As in the previous case, this is because taking into account the victim's loss of utility gives a more accurate estimate of the actual loss suffered by the victim and therefore it decreases the number of adverse commitments. When the loss has to be estimated, it can be overestimated or underestimated which is why the *Semi-Analytic Market Rule (with Estimated Opponent Effect)* performs worse than the policy that has full information about the loss. As in all other cases discussed so far, with extreme utility loss levels, the differences between the different policies are small, because either almost all contracts are decommitted from (high losses) or almost none of them are (low losses).

From these experiments, we can see that taking the possibility of opponent effect into account is therefore useful and the more accurate the information about that effect is, the better the outcome. All findings are consistent with hypothesis 2 and we can therefore accept it.

We will now discuss how compensatory policies perform against non-compensatory ones. We start by discussing the most complicated and also the clearest case where both parties can be affected and re-entries are allowed. Specifically, we can use figure 4.5 that was discussed in the previous subsection. Here, we can see that the *Optimal Market Rule* outperforms all non-compensatory policies at loss levels of $0.2 - 0.5$ (at $p < 0.001$ level) and in addition, with losses of 0.1 and 0.6 (at $p < 0.01$ level). Also the variations of the *Semi-Analytic Market Rule* will be able to outperform all the non-compensatory policies at least in one data point. For the basic *Semi-Analytic Market Rule*, this occurs at the loss levels of 0.3 (at $p < 0.05$ level) and 0.4 (at $p < 0.01$ level). When the information about the victim's losses are taken into account the difference becomes clearer; the *Semi-Analytic Market Rule (with Estimated Opponent Effect)* beats all non-compensatory policies at $0.3 - 0.4$ (at $p < 0.001$ level). And when the information about the opponent's

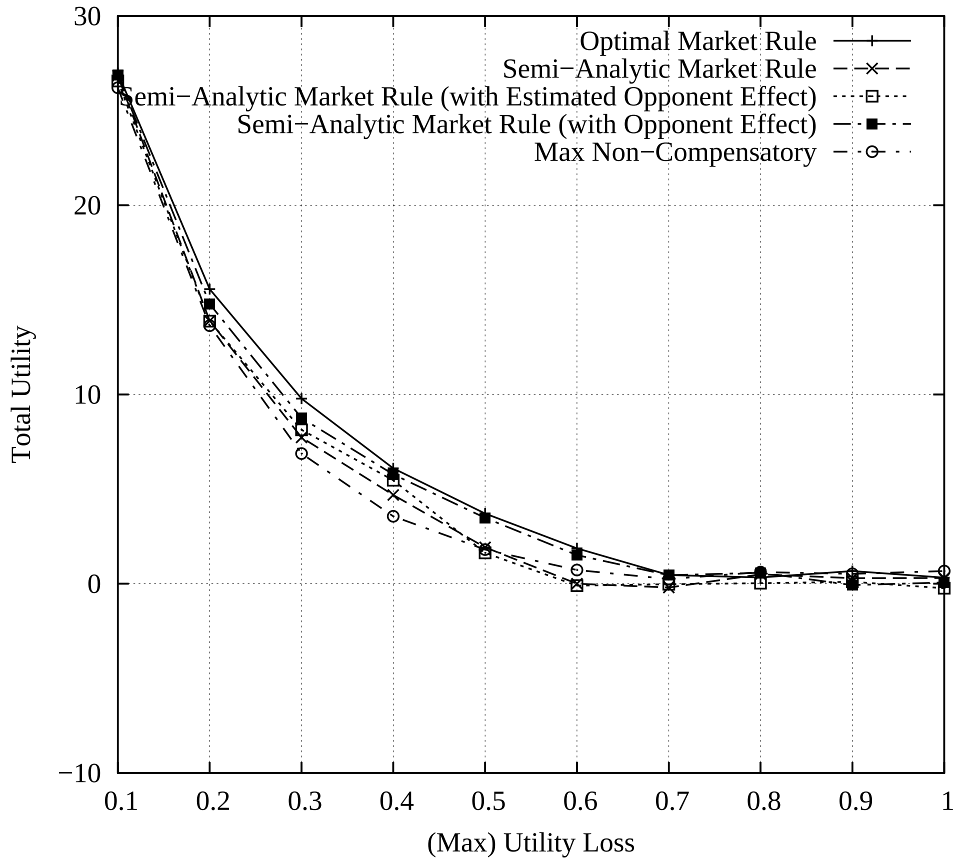

FIGURE 4.5: Effect of Opponent's Utility Losses (Hypotheses 2 and 3).

losses are complete, as in the *Semi-Analytic Market Rule (with Opponent Effect)*, the compensatory policy will outperform all non-compensatory policies at loss levels of 0.2−0.5 (at $p < 0.001$). Compensatory policies are therefore clearly better, although it is equally clear that their performance depends on the accuracy of the information they have. The better the information, the better the performance.

Also in the case where only the buyer can be affected but re-entries are possible (figure 4.6), we can see that the *Semi-Analytic Market Rule* outperforms the best of the non-compensatory policies between the loss levels of $0.2 - 0.5$ (at least at $p < 0.01$ level). In addition, the *Optimal Market Rule* policy is better than any non-compensatory policy with loss levels of $0.2 - 0.6$ (at $p < 0.001$ level).

Also in all the other cases, the optimal policies (Optimal Market Rule, Expectation Damages and Expectation Damages (with Opponent Effect)) are able to outperform the best of the non-compensatory policies at least in one loss level (at $p < 0.01$ level). As explained in section 4.2.1, this is because when the fees are equal to the actual losses of the victim, there are no adverse commitments

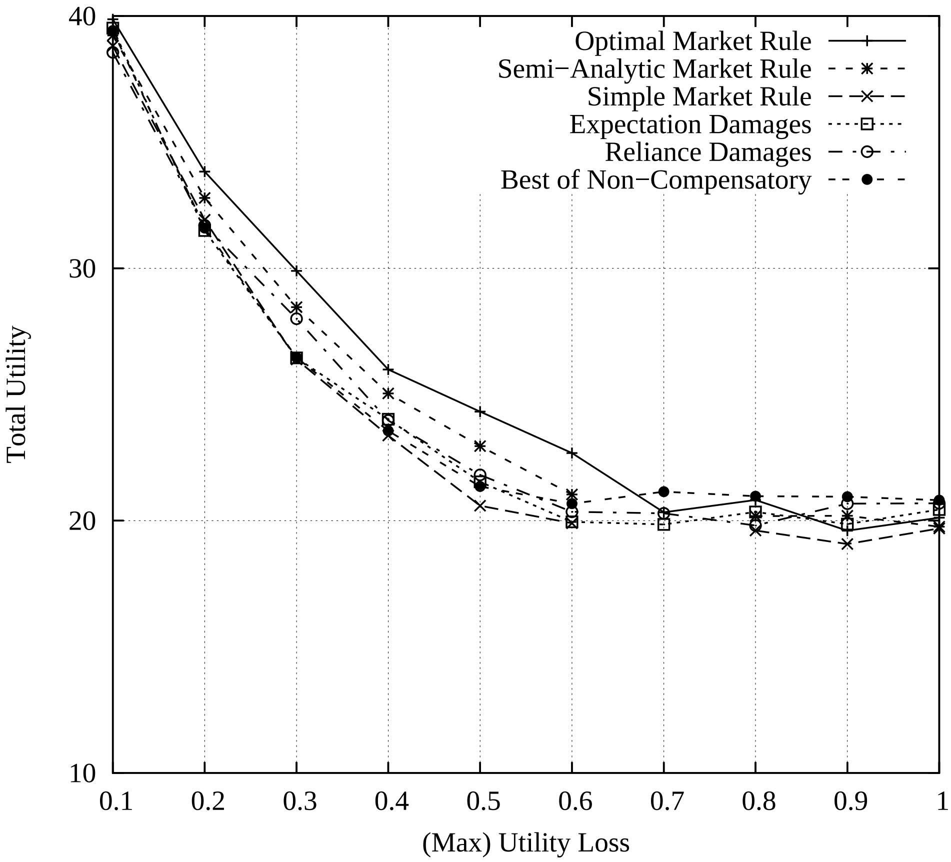

FIGURE 4.6: Semi-Analytic Market Rule vs. Best of Non-Compensatory when the buyer can be affected (Hypotheses 3 and 4).

and decommitments, but decommitments are always optimal (only occur when the contract has become detrimental for the common good).

However, incomplete information policies do not perform this well, because they cannot help causing some detrimental commitments and/or decommitments. In the case where no re-entries are allowed but both parties can be affected, only the *Analytic Compensatory (with Opponent Effect)* policy will be able to beat all non-compensatory policies at loss levels of 0.2 and 0.4 (at $p < 0.01$) and in cases where re-entries are not allowed and only one of the parties is affected, the *Analytic Compensatory* policies fail to outperform the best non-compensatory policies. This is because the *Last Half* deadline setting means that the estimates used by the *Analytic Compensatory* policy are not very accurate. For an explanation, see figure 4.7(a). Here, the loss for the seller is either $p - c$ (his profit) if he has not paid his costs or p (profit + cost) if he has. Now, the seller pays his cost at some point between turns 500 and 1000. Since all turns are equally likely the estimated loss (the fee) increases linearly from $p - c$ to p over this time. However, because the actual cost is paid at one point in time, this estimate overestimates the loss until the cost actually is paid and underestimates it after that until the time of performance. The error is shown as two triangles.

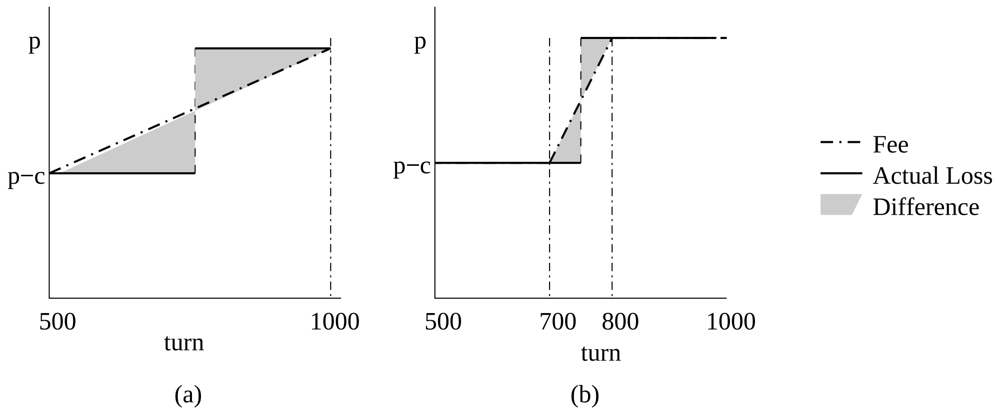

FIGURE 4.7: Analytic Compensatory policy estimates on the seller's loss.

The situation clearly changes when the deadline setting is changed to *Turns 700-800* (see figure 4.7(b)). In his case, the triangles showing the error are clearly smaller and therefore, on average, the estimates used by the *Analytic Compensatory* policy are more accurate. Now, in the case where both parties can be affected, even the basic *Analytic Compensatory* policy is able to beat the best non-compensatory policies at a loss level of 0.4 (at $p < 0.01$ level), *Analytic Compensatory (with Estimated Opponent Effect)* does the same more clearly (at $p < 0.001$ level) and *Analytic Compensatory (with Opponent Effect)* outperforms all the non-compensatory policies at loss levels between 0.2 and 0.5 (at least at the $p < 0.01$ level, in most cases at the $p < 0.001$ level). The results improve clearly on the case where the buyer is affected. Here the *Analytic Compensatory* policy is able to outperfom all the non-compensatory policies at the loss level of 0.4 (at $p < 0.001$ level) and also at loss levels of 0.3 and 0.5 (at $p < 0.05$ level). This improvement in performance occurs because the *Analytic Compensatory* policies have much more accurate information about the losses of the opponent and therefore the number of both adverse decommitments and commitments decreases.

Making it easier to estimate the providers' costs more accurately does little to help in the case in which the seller is affected, because the seller's costs do not affect the buyer's losses. However, we can still show that compensatory policies are useful also under incomplete information when the seller is affected. Specifically, when we set the effect probability to 1.0 (effect is certain), the *Analytic Compensatory* policy is able to outperform all non-compensatory policies at loss levels of $0.4-0.6$ (at $p < 0.001$ level). This means that the *Analytic Compensatory* policy is superior for the cases in which the effect occurs, but when the difference is so small that when there are many cases where utility losses do not occur, the difference may not be statistically significant. This is due to the fact that the value of a service can vary considerably in our set-up and the estimates the *Analytic Compensatory* policy uses are therefore not very good. For example, for quality 0.49, the value can be anything from 0 to 0.49 and for quality 0.99 it can be anything from 0.5 to 0.99. In a setting, in which we could estimate the value more accurately, the results would be better.

In summary, the compensatory policies can perform better than their non-compensatory counterparts, both under complete and incomplete information in all of our settings. This finding is consistent with hypotheses 3 and 4 and we can therefore accept them both.

Finally, we investigate the relationship between the average compensation error and the average total utility. As explained in section 4.2.1, we expect the average compensation error and the average total utility to be inversely related (hypothesis 5).

We ran the experiments as described earlier and for each decommitment for each run, we calculated the compensation error (the absolute value of the difference between the fee and the actual loss). For each experiment we then calculated the average compensation error and the total utility achieved. To investigate the effect on overall performance, we calculated the averages of both the total utility and the compensation error over all 10 loss levels ($0.1 - 1.0$). Due to our set-up and methodology, we could only calculate the compensation errors for the cases where decommitment actually occurred and, therefore, the cases where the fee did not allow any decommitments are not included. Out of 107 non-compensatory decommitment policies, 46 allowed commitments at all loss levels and are therefore included.

The results for the case where both parties can be affected and re-entries are allowed are in figure 4.8. Other cases are very similar and there is no reason to show them here. From the figure, it is clear that the average loss and compensation error are correlated. The correlation for this case is -0.89, which is different from zero in a statistically significant way (at $p < 0.001$ level). Also it can be noted that when the average compensation error becomes small (under 0.1 or so) the differences in terms of total utility are small. This is because these policies are able to handle most situations correctly and the difference occurs in a very small subset of cases. However, when the error is bigger, the difference is quite stark, because the bigger the error, the more cases are affected and adverse decommitments and commitments become more common.

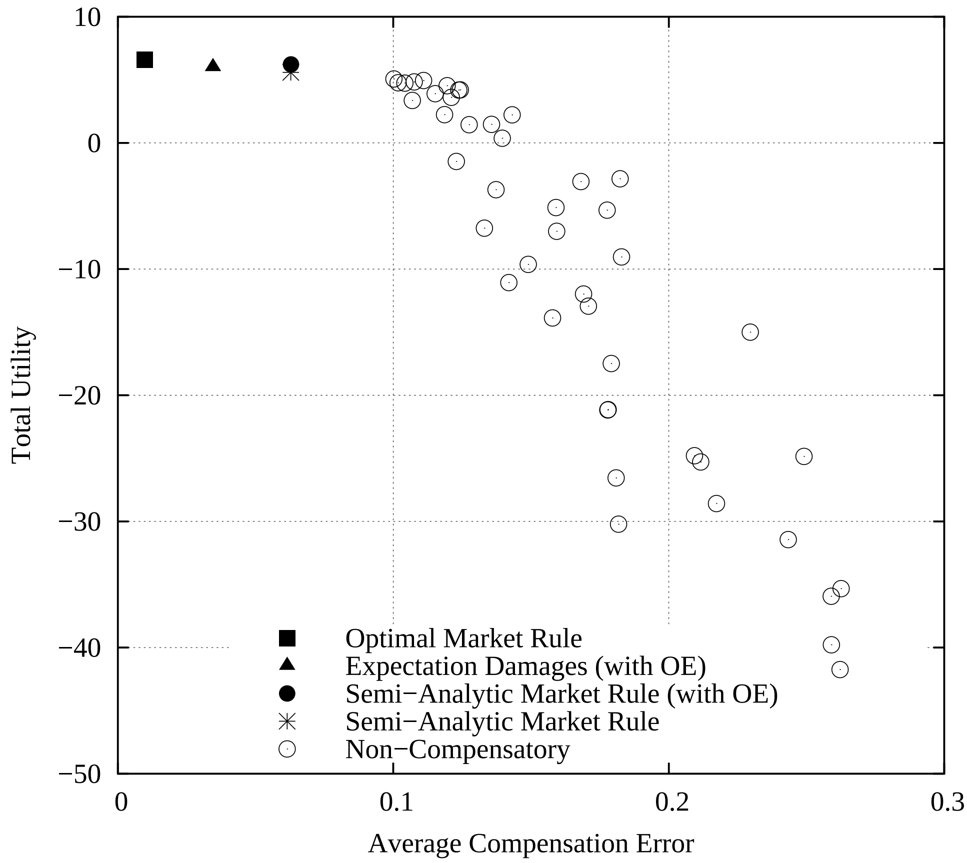

FIGURE 4.8: The effect of average compensation error on average total utility when re-entries are possible and the buyer can be affected (Hypothesis 5).

There is clearly an inverse relation between the average compensation error and total utility, which was what hypothesis 5 suggested. We therefore accept it.

4.3 Summary

In this chapter, we discussed performance decision in many different settings that included settings with one and two-sided decommitments, re-entries and incomplete information (our contribution **C1**). We showed that compensatory decommitment policies for contracts in electronic marketplaces can improve the welfare of the society in many settings and that ideas from contract law can be successfully used to provide useful new compensatory decommitment policies. It is also worth mentioning that the compensatory policies perform especially well in more complicated settings where both parties can be affected and/or re-entries are allowed.

This is because simple non-compensatory policies are much less accurate in such environments and, as we also showed, there is a clear inverse relation between the compensation accuracy and performance.

In more detail, we established that the compensatory policies have to be adapted to the environment in which they are to be used and they need reasonably accurate information about the losses to be useful. This means that they may not be suitable to all environments and situations, but where they are suitable they can be used to improve the total utility of the market participants. We showed that allowing parties to re-enter the market to find a substitute contract in the case of decommitment can improve the total utility, but only if the chances of finding a useful replacement (one that does not get abandoned later) are at least reasonable. Moreover, we showed that the opponent's possible utility loss should be taken into account where applicable.

We will now move on to the reliance decision, where a contract also already exists, but a party needs to decide how much to rely on the promised performance.

Chapter 5

The Reliance Decision

The *reliance decision* is about how much to rely on the promised performance. In our setting, it is more naturally a decision for the consumers, because for the sellers, the transaction is about money as is the compensation.[1] However, the buyers can get other services and use other resources to increase the value of the service the seller has promised to deliver and the level of this spending is a much more interesting problem.

To illustrate the issues in reliance decision, we therefore discuss a situation where the buyer has to decide how much to rely on the seller's promised performance. Specifically, the reliance means that the buyer takes (costly) preparatory actions that enhance the value of the seller's performance but which may turn out to be useless if the seller doesn't perform its part of the bargain. The more the buyer relies on the seller's performance, the higher these (potentially useless) costs and the higher the combined value of the seller's performance and these preparatory actions are. We assume that these reliance costs have to paid immediately after the contract in full and whether or not the service is later performed makes no difference to these costs. In case of non-performance, it is assumed (for simplicity's sake) that these costs are wasted and no benefit can be obtained from the preparatory actions without the actual service. The work here is associated with our research contribution **C2**.

We will first discuss the problem in more detail and explain how we have enhanced the basic model in order to investigate the reliance decision (section 5.1). We will

[1] Of course the sellers will have to pay their costs and if it is very likely that the buyer will not want the service in the end, the cost might not be worth spending if it is not covered by the decommitment policy.

then investigate these ideas empirically (section 5.2) and conclude by summarising our findings (section 5.3).

5.1 The Problem and the Modified Model

Unlike the basic model, we allow the consumer to increase the value of the service the provider is expected to deliver by acquiring complementary services or making other similar adjustments to his behaviour. The drawback of making these adjustments is of course that they take resources that could have been allocated to other tasks, they for example cost money. Thus, if the provider does not perform the service as agreed, these resources are wasted, because we assume that they offer only neglible value on their own and can only be used to enhance the service that never materialised. As more and more of these complementary services are purchased, both the value of the combined service and the cost of these additional services increase. We assume that the value of the service always increases more than the cost the buyer has to put in, because otherwise no seller would put in the extra effort. On the other hand, this means that if the buyer could be certain that the seller will perform its promised service, the buyer would therefore always rely on the performance fully. But when the performance is uncertain, the buyer may be better off getting only some of the complementary services or even none at all, if the performance is unlikely.

To make things simpler we introduce separate and discrete reliance levels. In particular, we limit the possible levels of reliance to five (see table 5.1). The lowest level (minimal) is the same as in the basic model, so the value is at most equal to q and the (reliance) cost for the buyer is equal to zero.[2] The other levels (limited, moderate, heavy and full in order of increasing investments) require more investments and deliver better value if the provider does indeed deliver the service. To investigate different types of situations, we have three settings described in the table below. Now, all settings have exactly the same characteristics at the lowest and highest reliance level, but they vary in between.[3] In more detail, in setting 1, the cost increases at a constant pace. In setting 2, the cost increases slowly at first, but faster and faster with every reliance level. And in setting 3, we have the reverse case, where the cost increases rapidly at low levels of reliance and more

[2]The value is exactly q when $q_b^{min} \leq q \leq q_b^{max}$, otherwise the value is lower.

[3]Also the value in the case of performance is always the same, $1q, 1.5q, 2q, 2.5q$ and $3q$ at different levels of reliance. Only the cost that has to be paid for the reliance varies and with that also the net value (=value−cost) of the success.

Reliance Level	Setting #1		Setting #2		Setting #3	
	$V_{b,1}(e,q)$	$C_{b,1}(e,q)$	$V_{b,2}(e,q)$	$C_{b,2}(e,q)$	$V_{b,3}(e,q)$	$C_{b,3}(e,q)$
Minimal	1q	0	1q	0q	1q	0q
Limited	1.25q	0.25q	1.4q	0.1q	1.1q	0.4q
Moderate	1.5q	0.5q	1.7q	0.3q	1.3q	0.7q
Heavy	1.75q	0.75q	1.9q	0.6q	1.6q	0.9q
Full	2q	1q	2q	1q	2q	1q

TABLE 5.1: Reliance Settings.

slowly in the high levels. In all cases the reliance cost depends on the quality of the original service reflecting the fact that higher quality services need also higher quality (more expensive) complementary services.

The sellers are heterogeneous both in quality (as in the basic model) and in reliability. The seller's reliability, ρ_s, denotes a probability that it will perform the service as agreed. The reliability is selected independently at random from the same distribution. We use four different well-known distributions: *Uniform*, *Normal*, *Exponential* and a derivative of the last one that we call *Reverse Exponential.* We use the *Normal* distribution with an average of 0.5 and a standard deviation of 0.2 and we have $\lambda = 5$ for the *Exponential* distribution.[4] The *Reverse Exponential* distribution is the distribution of $1 - X$, where $X \sim Exponential(5)$ and $X \in [0, 1]$. We choose these four distributions because they offer different settings where all reliabilities are equally likely (Uniform), where reliabilities around are 0.5 (Normal) and where reliabilities are more likely to be at the low end (Exponential) or at the high end (Reverse Exponential). To illustrate the distributions, consider the probability density functions in figure 5.1. When a seller enters a contract, a variable is selected at random from the interval $[0, 1]$ and if this variable is greater than the seller's reliability, it will have to decommit at some point before the delivery time. The decommitment time is selected at random from the distribution $Uniform(t_{contract} + 1, t_{delivery} - 1)$. This mechanism replaces the adverse effects used in the basic model. As with quality, the seller's reliance is assumed to be known in the market by all parties.

With the three reliance settings, these four reliability distributions give us a total of 12 settings. The decision-making for the buyer is exactly the same for all of

[4] Of course since both the Normal and Exponential distributions are continuous and unbounded and the reliability is bounded to the interval $[0, 1]$, we have made relatively straightforward changes. And quite simply, we use these distributions repeatedly until we get a value that is in the required interval $[0, 1]$. The probability of that occuring at the first attempt is 99.3% for the exponential distribution and 98.8% for the *Normal* distribution.

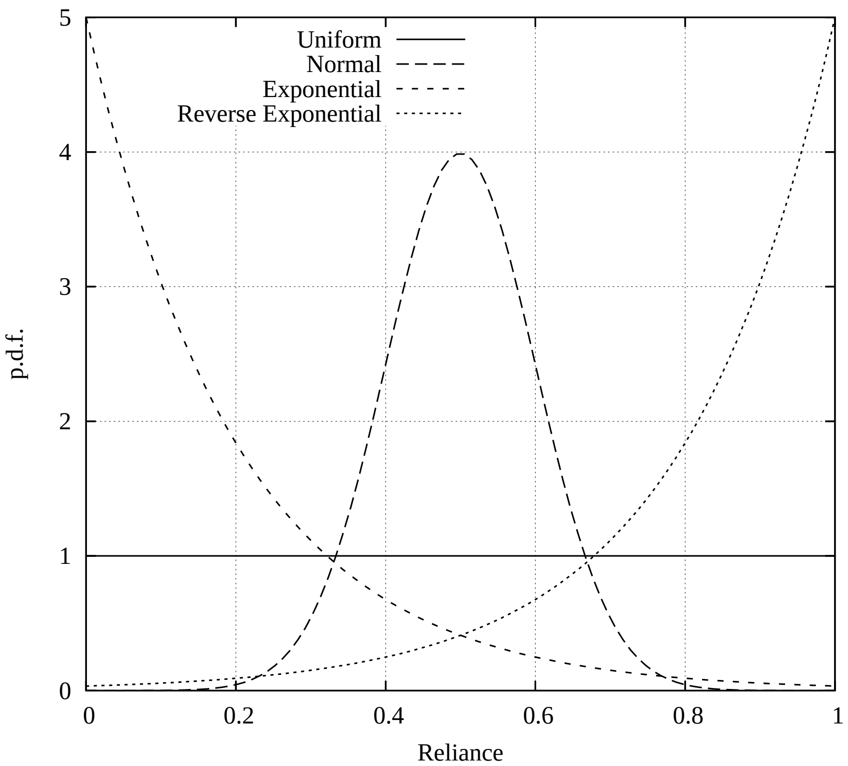

FIGURE 5.1: Reliability Distributions.

them. The buyer tries to maximise its expected utility. The utility for the buyer is:

$$U_b(p, e, q) = \rho_s V_{b,*}(e, q) + (1 - \rho_s)(f - C_{b,*}(e, q))$$

where e is the level of reliance and $*$ marks the setting. The reliance decision occurs after the contract has been entered into and we assume that it is *not* taken into account when the consumer is negotiating with the providers. The consumer will not therefore offer a higher price to ensure a deal with a highly reliable seller that provides excellent quality services, because that would take us to the realm of another decision, namely the contract decision that is to be discussed later, in section 6. So here we simply assume that the consumer gets a set of contracts (gained through negotiations as per the basic model) and decides his reliance given the contract and the characteristics (quality, reliability) of the service provider.

Of course an intelligent strategy is to choose the level of reliance so that it maximises the buyer's expected utility given the quality of the provider and the contract price. If the decommitment fee f does not depend on the reliance level (as is the case in all non-compensatory policies and also in partially compensatory policies that only compensate for the *minimal* reliance no matter what the real reliance level is), it will be the same no matter what reliance level the consumer chooses. Now, let us assume that the consumer has an existing contract (p and q are fixed) with a certain provider (ρ_s fixed) and it is making a choice between two reliance levels e_1 and e_2. The consumer b will choose the reliance level e_1 over e_2 if and only if:

$$U_b(p, e_1, q) \geq U_b(p, e_2, q)$$

If we put in the equation for $U_b(p, e, q)$ (above) and rearrange we get:

$$\rho_s(V_{b,*}(e_1, q) - V_{b,*}(e_2, q)) + (1 - \rho_s)(C_{b,*}(e_2, q) - C_{b,*}(e_1, q)) \quad \geq \quad 0$$

In other words, both p and f are irrelevant in the decision-making at this stage. A higher (but reliance-independent) decommitment fee does not therefore increase a consumer's reliance. A higher fee will obviously be better for the buyers, though, but this only changes the wealth distribution between the buyers and the sellers. The situation changes, if the decommitment fee depends on the reliance level, because instead of one f we would get two different fees (f_1 and f_2) and they stay in the equations. In other words, the decommitment policy does have an impact on the level of reliance, if it compensates for the reliance costs or expected reliance profits. Given this, we investigate the effect of two policies:

- *Expectation Damages (X% of Reliance Covered)*: The seller compensates for the buyer's profit and possible costs. In all cases, the buyer's profit with reliance level *Minimal* will be fully covered, but only X% of his reliance costs and extra profit with higher levels of reliance, where $X \in \{0, 25, 50, 75, 100\}$.

- *Reliance Damages (X% of Reliance Covered)*: The seller compensates for some of the buyer's possible costs. This means X% of the reliance costs, where $X \in \{0, 25, 50, 75, 100\}$.

It is easy to see that at $X = 100\%$ both policies encourage full reliance in all our settings. This is because the *Expectation Damages (100% of Reliance Covered)* ensure in the case of non-performance that the buyer obtains the same utility he would get in case the service is delivered and this utility is, in all three settings, at its highest (up to $2q$) when the reliance is full. In the case of *Reliance Damages (100% of Reliance Covered)*, the costs are always fully covered, so the buyer will always get zero profit in case of non-performance. However, the utility for performance is at its highest at *Full* reliance, so the expected utility will also be at its highest then. When no reliance costs are covered (0%), both policies make the parties choose the reliance level optimally (taking into account the losses in case of non-performance and, critically, the reliability of the provider).

In law, the problem of over-reliance is partially solved by compensating only for the losses (and profits) that are 'normal' for the service in question. This means that very unusual or unforeseeable (from the opponent's point of view) costs and expected profits are not compensated, even if real. If a party to a contract wants a wider protection to his interests, he needs to inform his opponent of his specific risks and possible losses before the contract is agreed on, so that they can be taken into account in price and other terms of the contract.[5] Different levels of reliance might be considered 'normal' in different situations and with different services. Here, however, we will investigate two policies:

- *Expectation Damages (Reliance Cover Capped at Y)*: The seller compensates for the buyer's profit and reliance costs, but only up to a reliance level $Y \in \{Minimal, Restricted, Moderate, Heavy, Full\}$. Any reliance above this 'normal' level will be completely uncovered.
- *Reliance Damages (Reliance Cover Capped at Y)*: The seller compensates for the buyer's reliance costs in full up to a reliance level $Y \in \{Minimal, Restricted, Moderate, Heavy, Full\}$, but any reliance above this 'normal' level will be completely uncovered.

In both of these policies, the cover will therefore increase to the maximum covered level and after that the compensation will only cover that maximum amount even if the actual reliance was higher. On the other hand, if the actual reliance was lower, the cover will be lower too. In our setting, the buyer will always select the reliability level that is at least the maximum covered by these policies. This follows from the same logic that was explained with the other two policies: the expected utility is always bigger with bigger reliance when the cover is complete. However, when it comes to reliances over this minimum level, the buyer will optimise his reliance appropriately. This policy will also lead to over-reliance and the higher the covered level, the higher the over-reliance will be.

We also use cases where the level of reliance is either selected at random or set as one of the five reliance levels in all contracts to illustrate the usefulness of the optimisation approach.

[5] Often parties also protect themselves by using liability limitation clauses in the contract. This means that a party's liability in case of non-performance is limited to a certain amount and if this is held to be reasonable limitation, the courts will not order damages that exceed this amount.

5.2 Empirical Evaluation

We will now investigate our ideas empirically. We explain what we did in three parts. First, we will discuss our hypotheses (section 5.2.1). We will then shortly explain how the experiments and statistical analysis to test the hypotheses were conducted (section 5.2.2) before, finally, discussing the results (section 5.2.3).

5.2.1 Hypotheses

There are a number of simple reliance strategies that the buyers are able to use. We investigate the following:

- *Random*: The consumer chooses a reliance level at random. All levels are equally likely.
- *Constant (Y)*: The consumer always chooses the same reliance level. There are five different variations, one for each level: $Y \in \{Minimal, Restricted, Moderate, Heavy, Full\}$.

None of these minimal policies should of course be able to beat our optimising policy that uses more information and more sophisticated mathematics in choosing the reliance level. However, sometimes the setting can be such that the optimal reliance policy is dominated by one of the reliance levels and then one of the simple *Constant* strategies may come to close in the terms of performance. Specifically, we contend:

> **Hypothesis 6.** The reliance optimisation strategy is always at least as good as any simple reliance strategy (random/constant) and if the optimal strategy is not dominated by one of the reliance levels, it will be better.

As discussed in the previous subsection, the full compensation for the consumer's reliance costs (and profits) will lead to full reliance always and with every contract. This may often mean that the consumer will over-rely on the service, especially if the reliance of the provider is low and he is unlikely to perform the task in the end. The more of the buyer's reliance is covered by the seller's decommitment policy, the more the buyer over-relies on the service. We therefore contend that:

> **Hypothesis 7.** In both *Expectation Damages (X% of Reliance Covered)* and *Reliance Damages (X% of Reliance Covered)*, the policies with higher X are never significantly better than those with lower X and $X = 100$ is always worse than $X = 0$, especially if the optimal reliance strategy is not dominantly *Full* (there is more than a handful of unreliable providers).

The other type of reliance-dependent decommitment policy has a threshold reliance level up to which the reliance is fully covered and after which there is no cover. As explained, this means that the buyers will rely at least to the threshold level and in many cases this means over-reliance. The higher the threshold level, the less worried the buyer becomes about over-reliance and, therefore, over-reliance becomes more and more common and the performance should accordingly suffer in most settings. Thus, we contend:

> **Hypothesis 8.** In both *Expectation Damages (Reliance Cover Capped at Y)* and *Reliance Damages (Reliance Cover Capped at Y)*, the policies with higher threshold reliance levels are never significantly better than those with lower levels. Moreover, the case where the cover is full (up to *Full*) is always worse than cases where the cover is only up to the *Minimal* reliance level, if the optimal reliance strategy is not dominantly *Full* (there is more than a handful unreliable providers).

5.2.2 Experimental Setup

We ran the market in all 12 settings (four provider reliability distributions and three reliability cost settings) with various decommitment policies (all variations of Expectation Damages and Reliance Damages (X% of Reliance Covered and Reliance Cover Capped at Y)) with optimised reliance strategy and various reliance strategies (Random and all variations of Constant) with *Constant 0.00* decommitment policy (this policy has no effect on the agent behaviour). The Expectation Damages (0% Reliance Covered) and Reliance Damages (0% Reliance Covered) were used as an example of reliance-independent decommitment policies. Since the policy does not have an effect on the optimal reliance level, no other reliance-independent policies were tried. In simple reliance strategies the decommitment policy does not matter (to the total utility), so no other policies were needed there either.

In each setting and decommitment policy or reliance strategy, we ran the simulation 100 times. We then did a one-sided t-test to the two averages compared.

5.2.3 Results

In figure 5.2, we have the two optimised cases compared to the *Random* and *Constant* (minimal, limited, moderate, heavy and full) reliance strategies. In the majority of the cases, it is clear that the two cases using the optimised reliance strategies produce superior results to the cases using other strategies (at $p < 0.0001$ level). However, there are a couple of exceptions. First of all, the cases where reliance is very high (*Reverse Exponential* distribution), one of the fixed reliance strategies, *Full*, is almost as good as the two cases using the optimised strategy. And actually in setting 3, the optimised policies are better, but only at the $p < 0.05$ level and in setting 1 at the $p < 0.01$ level. Another interesting case is setting 2 when the reliances come from the normal distribution. Here, the *Moderate* fixed reliance strategy is statistically indifferent from the optimised version.

In all these cases, the optimised strategy fails to make a clearer difference because the optimal strategy heavily uses a certain reliance level. In the negative exponential distribution cases, the provider reliability levels are often very high, so the optimal strategy is to rely considerably on the performance and, in many cases, this means *Full* reliance. In settings 1 and 3, the optimised strategy opts for the *Full* reliance in over 90% of the cases.[6] In a similar fashion, the normal distribution case in setting 2 has the optimal strategy taking *Moderate* reliance almost 70% of the time. This is simply because the mediocre reliability level prevalent in the normal distribution calls for *Moderate* reliance in this setting and what is even more remarkable is that in almost all cases (over 99%), the optimal reliance level is either *Moderate* or one of its neighbours (*Limited* or *Heavy*). This means that the differences in expectated utility between the optimal and *Moderate* reliance are very small even if they are not exactly the same and the total difference, is actually small enough to be statistically insignificant (the difference drowns in the noise). In the *Reverse Exponential* distribution cases (settings 1 and 3), the optimal reliance policy in the remaining 10% of cases is almost always *Minimal*, which produces quite different results from the *Full* reliance and therefore a relatively small number of cases is sufficient to show some difference. All our observations are consistent with hypothesis 6 and we can therefore accept it.

[6] In setting 2, the same proportion is slightly over 60% (which is why the optimised strategy has less trouble there).

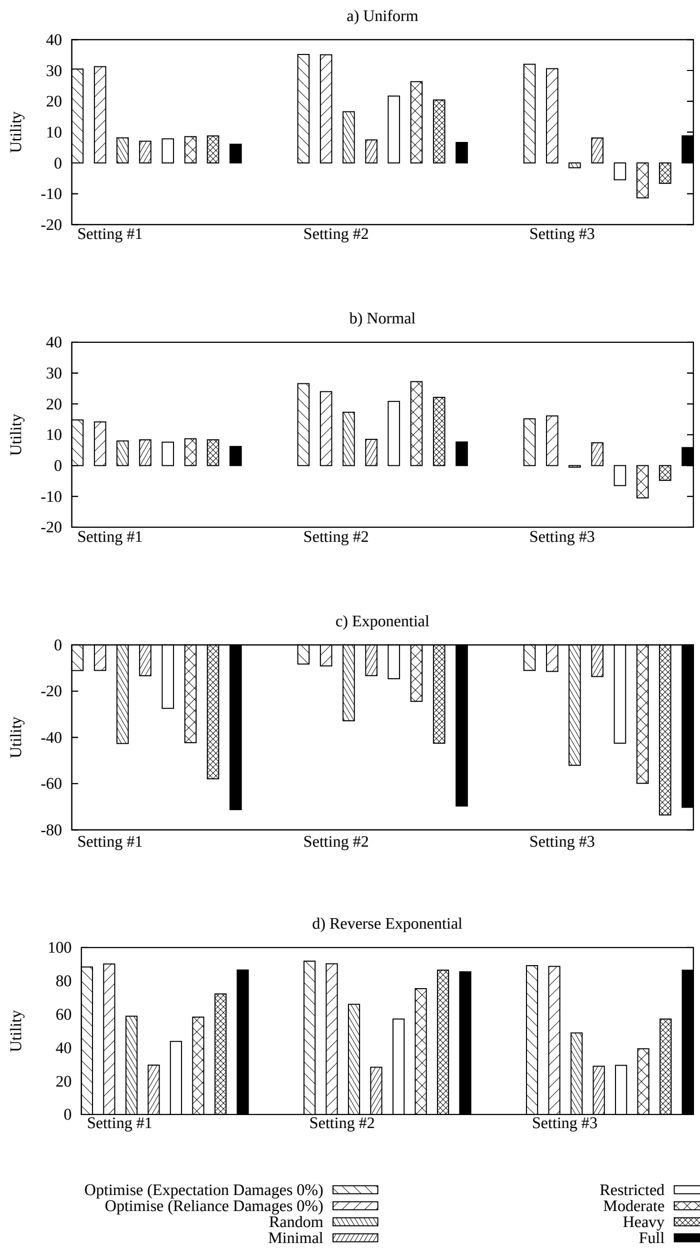

FIGURE 5.2: Optimised reliance vs. simple reliance (Hypothesis 6).

The second part was to investigate how compensating for the seller's reliance costs affects the utility. The results for all 12 settings are shown in figure 5.3. The downward trend when reliance compensation increases from zero to 100 percent is very clear in most cases and none of the occational increases are statistically significant (variances are quite significant). In the first three reliability distribution cases, both in *Expectation Damages* and *Reliance Damages*, the difference between no reliance compensation (0%) and full reliance compensation (100%) is statistically significant at the $p < 0.0001$ level. The differences in Reverse Exponential cases are, as expected, less clear. In most of the cases they are significant only at the $p < 0.01$ level and in one case, *Reliance Damages* in setting 3, no statistically significant difference can be found (it is very near though with $p = 0.0543$). The optimal strategy here is dominated by the *Full Reliance*, however, some 93% of the 176 average contracts use full reliance also in the optimal strategy. These are similar numbers to all cases in settings 1 and 3. In all cases, the average reliability in the market is around 0.8. We therefore accept hypothesis 7.

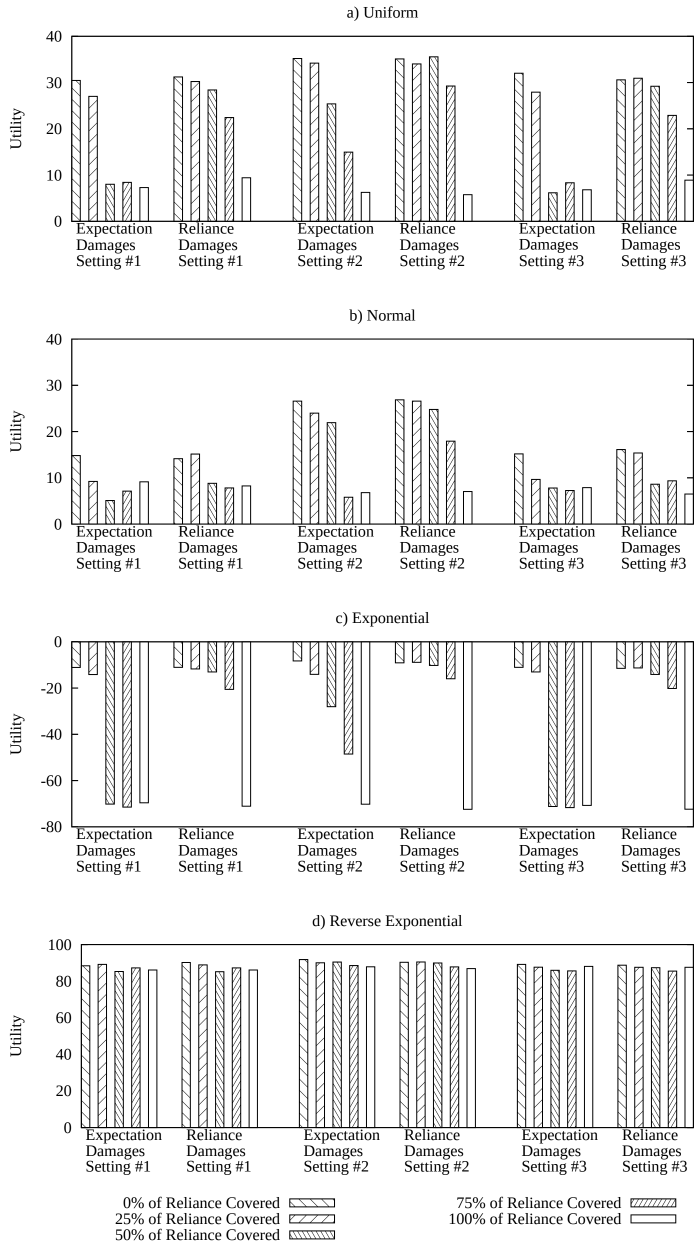

FIGURE 5.3: Effect of Partial Reliance Compensation (Hypothesis 7).

The final part has to do with capped reliance compensation, the results for which are shown in figure 5.4. As can be seen, when the compensation is restricted to the *Minimal* level, the situation is identical to the case where it was 0% and when compensation is full (up to *Full*) it will be equivalent to 100% reliance. Therefore it is hardly surprising that the results concerning the difference between minimal and full are similar to the previous case and we do not discuss them here again. The downward trend in the levels in between is even more clear than it was in the previous case. This is of course due to the fact that compensation up to a certain reliance level means that the buyer always relies at least to that level and, therefore, the increasing cover means that the buyer effectively has fewer reliance levels to choose from (because he has no incentive to select lower than the maximum covered reliance level). So, if an optimal policy with no reliance cover would use *Minimal* reliance level in some cases, under these policies, the buyer will always choose the lowest fully covered case and that means over-reliance from the society's point of view. There is no statistically signicifant increases and we can therefore accept hypothesis 8.

One additional observation from the last two figures is that when the reliances are very low on average (*Exponential* distribution), the total utility is negative even in the best cases. This of course is because most of the contract will not be performed, but a decommitment occurs at some point and that point may well be after the seller has paid its cost, making the total utility negative. The buyers will just get a set of contracts and will optimise their reliance on them. The optimal reliance strategy often calls for low levels of reliance in these cases and although that limits the damage considerably (compare the 0% and the 100% reliance cover cases when reliances come from the *Exponential* distribution), it is unable to stop it altogether. The one way to do that would be to let the sellers consider their reliability while negotiating on the price. In other words they would ask for a price that would keep the expected utility positive. In the next chapter, we will turn to these situations (although in a slightly different setting) when we investigate the contract decision.

5.3 Summary

In this chapter, we discussed the reliance decision and we showed that when the consumer is able to enhance the value of the service by making complementary

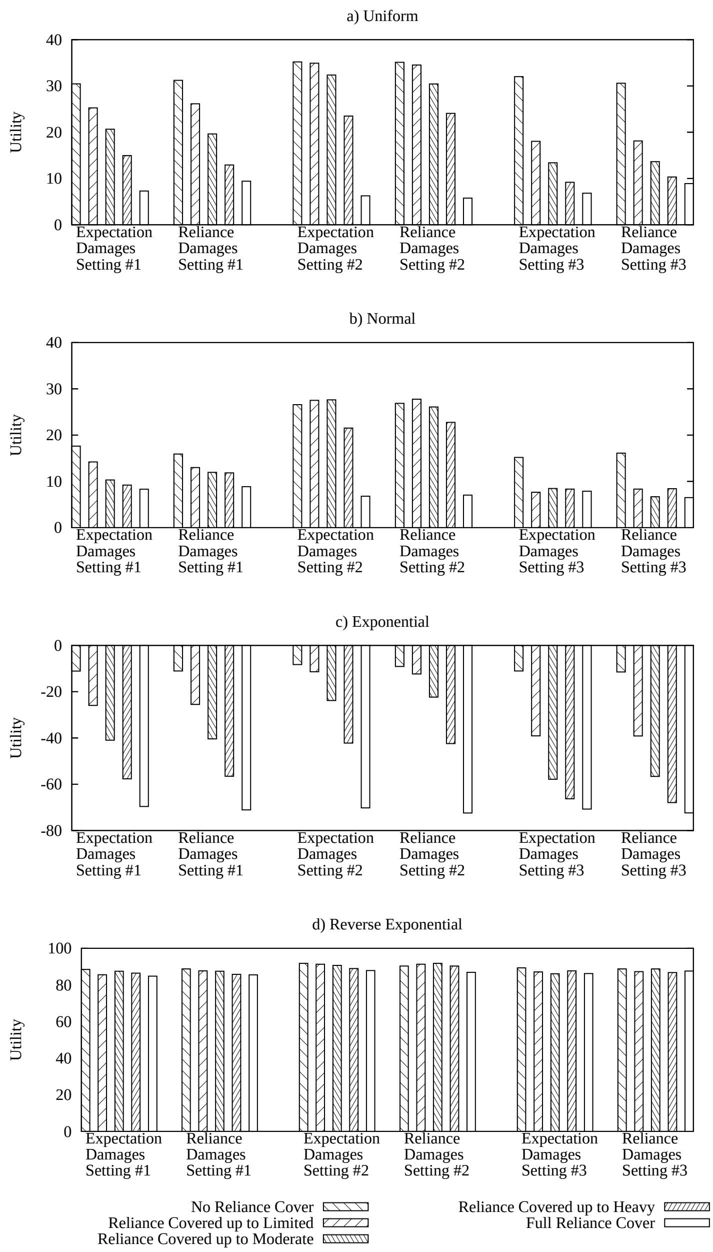

FIGURE 5.4: Effect of Limited Reliance Compensation (Hypothesis 8).

acquisitions, a decommitment policy that compensates for these extra costs (reliance costs) and/or profits will lead to over-reliance on the consumer's part. This over-reliance will lead to a decrease in total utility, the extra investments on the service that never comes are wasted. That is, they do not produce any benefit for the society, only costs. We also showed that limited or partial recovery of all reliance costs (and profits) is usually not a solution (contribution **C2**). A partial cover or cover that only comes to a certain point will decrease the amount of over-reliance, but it will not remove the problem completely.

From the society's point of view, the best policy is one that is *reliance-independent* (i.e. the fee does not depend on the level of reliance), so that it does not encourage over-reliance. Any such policy will be equally good in terms of the reliance decision, because no matter what the fee is, it will always be the same with all reliance levels and therefore won't affect the choice between these levels.

Chapter 6

The Contract Decision

In the first two decisions, the parties had a set of contracts and had to decide how to maximise their expected utility given these contracts. In the *contract decision*, however, the parties decide whether or not they should enter into those contracts in the first place. A contract can look beneficial, even lucrative, but if the circumstances can later change for the worse, it may turn into a disaster. As with the other decisions, the decommitment policies play an important role in the contract decision. Specifically, how an agent approaches a negotiation in a dynamic environment should clearly be very different, if the contracts can be decommitted from for free or if the decommitment fee is very high. And it is also crucial to know which of the parties might be affected in the future and how. A high decommitment fee is good news for an agent that is very reliable itself (needs to decommit very rarely), but whose opponent is very unreliable (likely to decommit), whereas a low fee would be beneficial in the opposite circumstances.

Here we will investigate how these concerns affect the behaviour of the agents and how that behaviour, in turn, affects the welfare of the society. This work relates to the research contribution **C3**. We will first discuss the issues in contract decision in more detail, as well as the changes we have made to our basic market model (section 6.1). We will then test our ideas empirically (section 6.2). Finally, we will conclude with the summary of the results (section 6.3).

6.1 The Problem and the Modified Model

The approach in this section is somewhat different from the others in this part of the thesis. Here, we do not aim to find the best decommitment policy and we do not use a large number of different policies. Instead, we investigate the basic forces in play. This is because otherwise it might be difficult to see what really is going on. For the same reason, our agents, when considering the contract decision, do not try to maximise their expected utility, but instead their aim is to achieve positive expected utility and any positive utility will do. By keeping the setting simple and by investigating only a few settings, we should be able to see the fundamental factors more clearly and we can then build more complex things on these basics in the next chapter, where the contract decision is one of the main problems we investigate.

The contract decision is different from the other decisions because it is much more flexible than the other decisions. This is because here the parties can use a contract price to re-distribute the risk between them. Here higher price means that the risk of getting negative utility in the end is lower for the seller and, similarly, a lower price means the lower risk for the buyer. This means that if the risk of negative utility is too high for one party (its expected utility is negative), the other party can, by offering a more advantageous price to that party and therefore by moving some of the risk to itself, encourage a party to enter in a contract. This means that the original liability rule, decommitment policy, is less important than in the other decisions. This can work in the society's favour in some cases, especially if both parties are taking the contract decision into account and parties can find a good balance between them by setting the price at suitable level. On the other hand, sometimes this can be problematic, especially in cases where only one party takes the contract decision into account and the other party will be able to transfer much of its risk to that other party without it even noticing. Sometimes this can be even detrimental to the common good.

However, this adaptation is also not always possible in full. This is more likely to be the case when the decommitment fee is very high and the contract price has set limits it cannot cross, for example if the price must be non-negative. We restrict the contract prices to the interval $[0, 1]$. This restriction has a practical purpose: Otherwise the prices could go very high or very low and it would be difficult to decide where to start the offers in a negotiation. On the other hand, it allows us to investigate the limits of the adaptation we just discussed. Even with such limits, the parties can do some adapting, the limitation just removes some of the more extreme versions of this behaviour (that again could hide the basic phenomena we seek to explore in this work).

We use four different decommitment policies:

- *Expectation Damages*: the decommitter compensates for the expected profit and the actual costs of the victim.
- *Reliance Damages*: the decommitter compensates for the costs of the victim.
- *Constant 0.0*: decommitment is free.
- *Constant 1.0*: decommitment is very costly.

These policies represent a wide spectrum of different types of decommitment policies. In the first one, *Expectation Damages*, the full cover means that the victim is not worried about decommitments. It is indifferent between performance and non-performance, because it is guaranteed the same utility in either case. Since it will only accept contracts with positive utility in the success case, it will get a positive utility in any such contract even in the case where the opponent decommits. The *Reliance Damages* case is similar. The cost coverage means that in cases where only one of the parties may decommit, the opponent is always guaranteed at least a zero profit, so it is unlikely to care too much about the risk of decommitment (given that any positive utility will do and if there is any possibility of success, the expected utility will always be positive). In these cases, the adaptation of the possible decommitter is therefore essential. The main difference between these two cases is the situation where both parties can be affected and are forced to decommit. The parties will then have to worry also about the possibility that they will have to cover for the opponent's profits and/or costs. However, only one party (the seller) has costs, so that will lead to some asymmetries in strategy.

In more detail, the *Constant 0.0* policy means that there is no compensation for the victim's loss. For the buyer, this is no problem, because it has no costs and therefore it will be guaranteed zero profit in case of decommitment, but the seller has to consider its possible costs. This means that the seller usually has to do the adapting in these cases, because the buyer will usually not. Whereas *Constant 0.0* often means under-compensation for the seller, the *Constant 1.0* policy will always mean over-compensation for both parties as victims. However, for the (potential) decommitters, it will mean trouble for both of them, because the fee will certainly make their utility seriously negative in case they have to decommit. Now, since they want to have their expected utility positive, this means that they need to ask for a higher (the sellers) or lower (the buyers) price. This case is interesting, because it forces both parties to take the fee into account and because it also makes the potential victim actually prefer non-performance (the utility of non-performance is higher than the utility of performance for the potential victim), so it may have an incentive to encourage the potential decommitters to take on contracts they would not otherwise do by giving them a higher or lower price. However, because the price is limited to the interval $[0, 1]$ and because the fee is so high, the adaptation using the price may not always be possible in full.

Because we are interested in whether or not the parties should enter a contract at all and because our evaluation would otherwise be mixed with the performance decision, we use a variation of the model where the negative impact is always so big that it will mean that the parties decommit. We also remove the limitation for the value that the buyers can get. The value can therefore be very negative, also for the buyers.[1] This means that the affected parties always decommit no matter what their decommitment policy is and, therefore, decommitment policies will have no effect after the contract is entered into. The effect will therefore all be whether or not the policy and its effect on the parties' decision making before entering a contract will improve the common good. In this context, it is obvious that if the parties take their own good (contract decision) into consideration, and avoid contracts that lead to negative expected utility, their utility will remain non-negative also when a probability of diastrous adverse effects is very high. When the effect probability is very high, the best course of action may be not to enter contracts at all, but take a zero profit instead. Otherwise the utility will be positive in expectation. In a single case, even careful parties may sometimes get a negative utility, because they want to avoid negative utility on expectation, not entirely, but in the long run and over many experiments, the average utility should always be non-negative. The more interesting question, however, is whether or not, this consideration can improve the common good and could this adaptation through price even be detrimental to the common good in some cases.

In the basic model, the possible effect always takes place between the contract and delivery times and all times in that range are equally likely (*Uniform* distribution). However, the problem with that approach here would be that the parties do not get new information about the possibility of the adverse effect before they have to decide whether or not to enter a contract. Therefore, the risk of decommitment would remain the same through the experiment. To facilitate improving

[1]In the basic model, the buyers' value was always at least -0.05 (see section 3.3).

information, we use an alternative pattern for the effects, namely *Uniform from Start*. This means that the adverse effect (that decreases the utility of any possible contract) can occur at any time during the experiment (between t_0 and $t_{delivery}$) and all possible times are equally likely. The effect occurs independently of any contract, so if a party is affected but has no contract, he will exit the market and if he is negotiating with somebody, he will withdraw from that negotiation immediately. If he is in a contract, he will decommit. However, if a party is still in the market at turn t, this must mean that he is either not going to be affected at all or he just hasn't been affected yet. The probability for the latter decreases as the time progresses:

$$P(\text{no effect so far} \mid \text{effect}) = \frac{t_{delivery} - t}{t_{delivery}}.$$

Now, to calculate the probability that the effect is going to take place later, we get (by Bayes' theorem):

$$\begin{aligned}
&P(\text{effect} \mid \text{no effect so far}) \\
&= \frac{P(\text{no effect so far} \mid \text{effect})P(\text{effect})}{P(\text{no effect so far})} \\
&= \frac{P(\text{no effect so far} \mid \text{effect})P(\text{effect})}{P(\text{no effect so far} \mid \text{effect})P(\text{effect})+P(\text{no effect at all})} \\
&= \frac{\frac{t_{delivery}-t}{t_{delivery}}a_{base}}{\frac{t_{delivery}-t}{t_{delivery}}a_{base}+(1-a_{base})}
\end{aligned}$$

where a_{base} is the effect probability at turn 0. So when $a_{base} = 0$, $P(\text{effect} \mid \text{no effect so far}) = 0$ and if $a_{base} = 1$, $P(\text{effect} \mid \text{no effect so far}) = 1 \forall t < t_{delivery}$.[2] Otherwise, the probability is decreasing in t as can be seen from figure 6.1. This means that every turn that a party stays on and does not withdraw, exit or decommit, the probability that he will be affected later decreases. At delivery time ($t = t_{delivery}$), the probability is always zero (when $a_{base} < 1$). So the later the parties enter into a contract, the better are the chances that the contract will actually be performed as agreed. However, if they wait too long, they might not have negotiation partners or their own deadline might arrive.

If only one of the parties can be affected, there are only two possible outcomes. Either the affected party will decommit later leaving the other party utility $U(\text{decommitment})$ or he will not (utility $U(\text{contract})$). The calculation for expected utility is therefore

[2] Obviously if $a_{base} = 1$ and $t = t_{delivery}$, this would be $0/0$ which is undefined, but that will never happen, because if the effect takes place certainly ($a_{base} = 1$), it will do so before the delivery time.

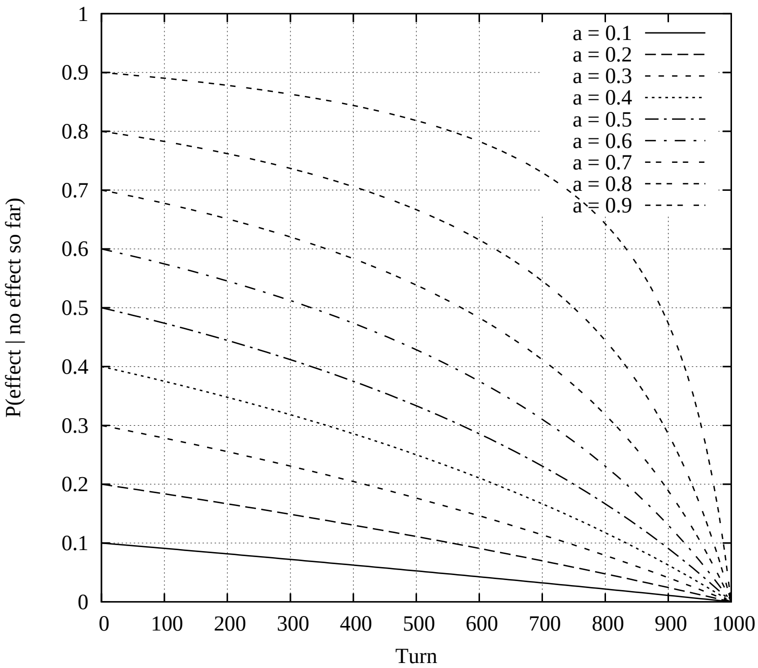

FIGURE 6.1: The effect probabilities over time $a \in [0.1, 0.9]$.

straight-forward:

$$\begin{aligned} EU = \ & (1 - P(\text{effect} \mid \text{no effect so far}))U(\text{contract}) \\ & + P(\text{effect} \mid \text{no effect so far})U(\text{decommitment}) \end{aligned}$$

Of course if both parties can be affected, there are four possible outcomes: neither, one or both parties are forced to decommit at some point. Accordingly:

$$\begin{aligned} EU = \ & P(\text{success})U(\text{contract}) \\ & + P(\text{buyer decommits})U(\text{buyer decommits}) \\ & + P(\text{seller decommits})U(\text{seller decommits}) \\ & + P(\text{both can decommit})U(\text{both can decommit}) \end{aligned}$$

where the probabilities are

- $P(\text{success}) = (1 - P(\text{buyer effect}))(1 - P(\text{seller effect}))$
- $P(\text{buyer decommits}) = P(\text{buyer effect})(1 - P(\text{seller effect}))$
- $P(\text{seller decommits}) = (1 - P(\text{buyer effect}))P(\text{seller effect})$
- $P(\text{both may decommit}) = P(\text{buyer effect})P(\text{seller effect})$

The only case worth further discussion is the one where both can decommit. In theory, the parties could decommit at exactly the same time and therefore would either get or pay no fees or would get and pay the fees as allocated by the applicable decommitment policy. However, in our case, decommitments cannot occur at exactly the same time because we only allow one party to a contract to decommit at any given time.[3] So in our setting, one party always decommits before the other. Since the effect at a random time for both parties (the same distribution) and parties do not try to outwait each other but decommit as soon as the effect occurs,[4] the parties consider it equally likely that they will in this case decommit first (and pay the decommitment fee) and be affected later (and be paid the decommitment fee). Should concurrent decommitment be possible, the parties would of course just estimate the probability of that and factor it in their calculation. For us, it would just add an unnecessary complication and, as mentioned, the possibility is omitted.

As explained, our agents will take any contract that gives them positive expected utility. They ensure positive expected utility by setting their reservation price so that this goal is always fulfilled if an agreement on price is reached. The reservation price is calculated again for each negotiator on each turn of the negotiation. We will now shortly discuss how the agents calculate the reservation price. This is reasonably straight-forward. We use the equation for EU above and write it on an inequality where we require it to be positive. In the case of a constant decommitment policy, we get for the seller:

$$\begin{aligned} EU_s(q,p) &= P(\text{success})(p - C_s(q)) + P(\text{buyer decommits})f + \\ &\quad P(\text{seller decommits})(-f - Ec(q)) + \\ &\quad P(\text{both may decommit})(\tfrac{1}{2}(f - f - Ec(q))) \\ &\geq U_{target} \end{aligned}$$

where U_{target} is the minimum acceptable utility (we used 0.00001 in our experiments). We then solve this inequation in relation to p and get:

$$\begin{aligned} p \geq\ & C_s(q) + \tfrac{1}{P(\text{success})}(U_{target} - P(\text{buyer decommits})f - \\ & P(\text{seller decommits})(-f - Ec(q)) - \\ & P(\text{both may decommit})(\tfrac{1}{2}(-Ec(q))) \end{aligned}$$

and this, of course, is the reservation price for the seller. It can be simplified in different ways in different settings. For example, if only the buyer can be affected we

[3] Both parties get their turn to decommit at every turn as if they were still negotiating.
[4] Such strategic decommitment was also ignored in the performance decision (chapter 4).

get $P(\text{buyer decommits}) = P(\text{both may decommit}) = 0$ and even if both can decommit, we investigate settings where $P(\text{buyer decommits}) = P(\text{seller decommits})$ and so on. The compensatory policies work in a similar way. For example, the reservation price for the buyer when the *Reliance Damages* policy is used becomes:

$$
\begin{aligned}
EU_b(q,p) &= P(\text{success})(V_b(q) - p) + P(\text{buyer decommits})(-Ec(q)) + \\
&\quad P(\text{seller decommits})(0) + \\
&\quad P(\text{both may decommit})(\tfrac{1}{2}(-Ec(q)) \\
&\geq U_{target}
\end{aligned}
$$

and for the seller when the *Expectation Damages* policy is used:

$$
\begin{aligned}
EU_s(q,p) &= P(\text{success})(p - C_s(q)) + P(\text{buyer decommits})(p - C_s(q)) + \\
&\quad P(\text{seller decommits})(-(V_b(q) - p) - Ec) + \\
&\quad P(\text{both may decommit})(p - \tfrac{1}{2}(Ec + V_b(q) + C_s(q))) \\
&\geq U_{target}
\end{aligned}
$$

The other cases go in a similar way. We are now ready to start discussing our results. The restriction is enforced simply by not allowing negative reservation prices for the seller (any negative reservation price is changed to zero) or reservation prices greater than one for the buyer (any such reservation price is changed to one). This means that if the seller's reservation price is above 1 or the buyer's reservation price is below 0, no contract can be achieved.

6.2 Empirical Evaluation

We will now discuss our empirical results. First, we will present our hypotheses (section 6.2.1). We will then explain how the experiments and the statistical analysis were conducted (section 6.2.2) before we finally, discuss the results (section 6.2.3).

6.2.1 Hypotheses

As always, only the performance of a task can increase the total utility (by $V_b(q) - C_s(q)$), non-performance leads only to the payment of decommitment fee (which does not increase welfare in the society, it merely moves it from one party to another) and possibly also the seller's costs ($Ec = D(t)C_s(q)$). The latter is a waste for the society, because it produces no benefit. So, the problem is to maximise the performances and to minimise the decommitments. Of course these two goals are contradictory: the decommitments would be minimised by having no contracts (no risk of decommitments), but the maximisation of performances might lead to a large number of contracts. With no contracts, the total utility in the market is zero. Taking some contracts, the utility might be positive and taking too many contracts, the total utility may even be negative (the costs of decommitments outweighing the benefits from the performances).

The basic idea is then to find a balance between the two and instead of getting just some contracts, getting contracts that are more likely to be performed. In our setting, where the buyers and sellers are homogeneous in terms of their reliability, the only way to improve the chances of success is to wait longer before entering into a contract. However, as explained, there are risks involved in waiting: if agents wait too long, they might not get contracts at all. Therefore the risk of decommitment has to be considered with a chance of getting a contract in the first place. When the risk of later decommitment is small, it may be more important to secure a contract whenever an opportunity arises. But when a decommitment is very likely, waiting is often more prudent.

As explained, the parties can use the contract price very effectively as a means of ensuring positive expected utility for themselves no matter what the decommitment policy is (within reason[5]) and if both parties are doing that then the expected utility for the society is always non-negative (positive if there are any contracts) in any setting in the long run. Sometimes it is even possible that one party alone, by taking the contract decision into account, can ensure positive expected utility. This requires that the only source of negative total utility, the seller's possible costs, is clearly in one party's responsibility and that party takes the contract decision into account. Other cases will not adapt correctly and will perform badly when the effect probability is high, because they will enter into contracts that are never going to be performed causing costs to the sellers. Thus, we contend:

> **Hypothesis 9.** When only one of the parties can be affected and the decommitment fee is either compensatory or *Constant 0.0*, the case where both parties take contract decision into account is no worse than any other policy and it will be better than the case where neither

[5]For example, *Not Allowed* or policies with *very* high decommitment fees, strict limits on the price or similar factors, can make this adjustation impossible and mean that society suffers (i.e. no mutually agreeable contracts can be found, although with less strict commitments such contract could be found).

> party considers the contract decision or the case where only one party does so, if the party in question is not liable for the provider's possible costs according to the decommitment policy. The difference is clearest with high effect probabilities.

As also explained earlier, serious over-compensation (like *Constant 1.0*) will make the decommitment on one hand, a potentially disastrous situation for the potential decommitter and, on the other hand, very attractive for his opponent. This means that the affected party that takes the contract decision into account will need to be very careful. Specifically, it will want to enter a contract only close to the deadline (the risk of decommitment is smaller) and/or have a very high utility (high reservation price for the sellers and low reservation price for the buyers) in the case it is able to perform to balance things out. In contrast, his opponent may well be interested in accepting very bad contracts (from its point of view), if that gives him a chance of getting the high decommitment fee if the opponent fails. On their own these two tendencies influence the decisions into opposite directions: when the affected party alone takes the contract decision into account, it will want to wait for longer and longer and this leads to less and less contracts when the effect probabilities increase until with very high effect probabilities there is (almost) no contracts. The very high fee makes this effect stronger and makes the potentially affected party go to great lengths to avoid too big a risk of that fee. So when the fee over-compensates the victim's loss as badly as the *Constant 1.0* policy does, the potentially affected party overdoes this (is too careful to avoid contracts) and this will affect adversely the total utility when it alone will consider the contract decision. The total utility is going to be worse than in the case where neither party takes the contract decision into account in intermediate effect probabilities. In very low probabilities there are not enough decommitments to make a significant difference and since staying out of risky contracts will ensure non-negative expected total utility it will be able to outperfom the case where neither will take the contract decision into account when the effect probabilities are very high and the expected total utility becomes negative.

On the other hand, when only the decommitter's opponent takes the contract decision into account, the effect is opposite: the opponent will want a contract as soon as possible. This is because it prefers the non-performance and is willing even to take a worse contract to secure a chance to get the high fee. The earlier it manages to lure the potential decommitter into a contract, the more likely the decommitment will occur. From the total utility's point of view, this is of course counter-productive and, therefore, in the case where only the victim takes the contract decision into account, the results are likely to be even worse than when no party takes the contract decision into account. This will happen especially with high effect probabilities when the fee is very likely.

In case the potential decommitter is the seller, there is also its own costs to consider. In case of decommitment, the seller will have to pay not only an overcompensatory fee but also its own possible costs. This makes the seller even more reluctant to enter into contracts and the effects described above are likely to be even stronger. Specifically, we contend:

> **Hypothesis 10.** When only one of the parties can be affected and the decommitment policy is *Constant 1.0*, both cases where only one of the parties takes the contract decision into account will perform worse than the case where neither takes the contract decision into account at least with some effect probabilities. For the case where the decommitter considers the contract decision, this occurs with low to intermediate effect probabilities and for the case where the victim considers it, especially with high probabilities. When the seller is the decommitter, these effects are stronger.

When only one of the parties is affected, the decision-making of a participant involves either the possibility that they themselves will have to decommit or that their opponent might have to decommit at some point. When both parties can be affected, both of these factors need to be considered at the same time. We only consider cases where the effect probability is the same for both parties. This means that the probabilities for the player itself or its opponent having to decommit is the same and in both *Constant* policy cases, also the decommitment fee is the same so these things cancel each other out in the case of the buyer and it is not therefore too worried about possible decommitments. However, for the seller the situation is significantly different because in addition to the fees, it has to consider the possibility that it will have to cover its own costs. This can occur both when it has to decommit itself and when the other party decommits (fees it pays or receives cancel each other out). This will mean that the seller will do all the necessary adjustments and the buyer none.

The compensatory cases are more interesting because the parties' profits and costs are different. Therefore the fees do not cancel each other out, instead they have to be considered. This means the parties will have to consider their profits and the costs they would receive as compensation if the other party decommits and their

Case	U_b	U_s	U_{b+s}
Buyer Decommits	$-f = -Ec_s$	$f - Ec_s = 0$	$-Ec_s$
Seller Decommits	$f = 0$	$f - Ec_s = -Ec_s$	$-Ec_s$

TABLE 6.1: Utilities in the *Reliance Damages* policy.

opponent's profits and costs that would have to be compensated for if they are forced to decommit. In the *Reliance Damages* policy case, either party considering the contract decision alone will not be able to cover all cases. This is because both parties alone consider only one of two cases. The buyer considers the utilities in the U_b column and the seller the utilities in the U_s column in table 6.1. Both of them ignore the seller's costs in one of the two cases, because either the buyer compensates them or ignores them. Only when both parties consider the contract decision at the same time will the seller's costs be considered by one party in all cases and therefore, the case where both consider the contract decision is likely to outperform other cases with intermediate effect probabilities (when there are enough decommitments, but there are still enough contracts).

The *Expectation Damages* case is a bit different. The expected utilities for each case are as shown in table 6.2. Now, under this policy, the buyer will get its full utility with the probability of $P(\text{success})+P(\text{seller decommits})+\frac{1}{2}P(\text{both may decommit})$ and it will have to pay the seller's profits and costs with the probability of $P(\text{buyer decommits}) + \frac{1}{2}P(\text{both may decommit})$. If the buyer's utility in case of success is high enough, compared to the compensation the buyer has to pay in case of its own decommitment, the buyer is willing to negotiate and enter a contract even when performance is very unlikely. In the extreme case, let $P(\text{buyer effect}) = P(\text{seller effect}) = 1$. This means that $P(\text{both may decommit}) = 1$ and all other probabilities are zero. Now, the buyer's expected utility is:

$$U_b = \frac{1}{2}(V_b(q) - p) + \frac{1}{2}(-p + C(q) - Ec_s).$$

This will be positive, if $V_b(q) - p > -p + C(q) - Ec_s$, so even if the decommitment is a certainty, the buyer is willing to enter a contract if the utility in case the seller decommits first $(= U(\mathit{success}))$ is greater than the fee it has to pay in case it has to decommit first. In other words, if the buyer is able to get a good contract (low price), it will be willing to take it even if it will know that it will never be performed. A similar logic applies to the buyer and the logic can also be extended to less extreme cases. Therefore when either party alone considers the contract decision, it can enter into contracts even if performance is very unlikely or even

Case	U_b	U_s	U_{b+s}
Buyer Decommits	$-f = -p + C(q) - Ec_s$	$f - Ec_s = p - C(q)$	$-Ec_s$
Seller Decommits	$f = V_b(q) - p$	$f - Ec_s = p - V_b(q) - Ec_s$	$-Ec_s$

TABLE 6.2: Utilities in the *Expectation Damages* policy.

impossible, as long as the utility in case of success is large enough compared to the fee it has to pay in case of failure. When the other party does not consider the contract decision, it will be willing to enter into such contracts. But when both parties take the contract decision into account, no contract can be acceptable to both of them when $P(\text{both may decommit}) = 1$ (for any given contract, both parties cannot have higher utility in case of success than in case of failure at the same time). Also in less extreme situations, the parties will be more careful. We therefore contend:

> **Hypothesis 11.** When both can be affected and decommitment policy is compensatory (either *Reliance or Expectation Damages*), the case where both parties consider the contract decision will always be at least as good as any other setting and it will be better at least some of the time.

The price can be effectively used between the parties to distribute the risks in an effective way. In this distribution, it makes little difference what the decommitment policy is, but all policies should perform roughly the same. However, given that the price is limited to the interval $[0, 1]$ it may well be that the parties are unable to adapt properly in case the fee is heavily overcompensatory and only one of the parties can be affected. This is because the party not affected is unable to set the price so that the party affected would be able to find that acceptable and not all parties will be able to wait until the risk would be acceptable (because of their deadlines and effects). The problem is less pronounced when both parties can be affected because the risk of an oversized decommitment fee is more evenly distributed. Thus:

> **Hypothesis 12.** When only one of the parties can be affected, the *Constant 1.0* policy will perform worse than the other policies.

The performance of different policies depends of course on the information they have. We have assumed that both the buyer and the seller will know when the seller pays its cost and therefore can calculate the expected cost of the seller accurately. However, already a small change like not knowing the time accurately will affect the buyer's ability to estimate his risks in the compensatory policies. So, if instead of $Ec = \frac{t_{delivery} - t_{cost}}{t_{delivery} - t_{contract}} C_s(q)$, all times after the contract is formed are equally likely and, then $Ec = \frac{1}{2} C_s(q)$. This sometimes overestimates and sometimes underestimates the seller's expected cost and will mean that the buyer sets his reservation price on a slightly wrong level. This small error alone is enough to decrease the total utility in cases where the buyer is affected and the buyer takes the contract decision into account. We therefore contend:

> **Hypothesis 13.** When the buyer is affected, but does not know the time the seller has to pay its cost, the performance in terms of total utility suffers in compensatory policies.

6.2.2 Experimental Setup

We ran the marketplace in 36 settings for each of the four decommitment policies. We used $a = \{0.0, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 0.995, 1.0\}$, for buyers and sellers alone (the other party had $a = 0$) and then for both of them at the same time. We needed to add the case $a = 0.995$, because in some situations a party's behaviour would change radically when the decommitment probability was 1.0. This is because our agents were after positive utility and if a decommitment produces zero utility and a performance positive utility then as long as there is any chance of performance a party will try to get contracts, but when the chances of performance drop to zero, the expected utility of any contract drops to zero and this is no longer positive so the party no longer enters into any contracts. The minimum positive value for utility is 0.00001, so accordingly the probability does not need to be exactly one, only very close to it. By taking an additional value very close to one, our graphs show more clearly that the difference in behaviour occurs very close to 1 and not gradually from 0.9. We could have of course changed the rule so that any non-negative utility would have been good enough and we would not get a break in continuity, but, on the other hand, entering into a contract is always a risky thing to do in our setting and there should be at least some benefit for doing so.

Again we ran the marketplace 100 times in each setting, calculated the averages and variances of various variables and did a one-tailed t-test to test whether or not the perceived differences were statistically significant.

6.2.3 Results

In the first hypothesis in this chapter, we claimed that when only one party can be affected and the decommitment policy is either compensatory or *Constant 0*, the total utility will be at least as high in cases where both parties take the contract decision into account than in other cases and that it will be higher than the cases where neither or only the party that is not responsible for the seller's possible costs takes the contract decision into account with high effect probabilities. We have the total utilities for these cases in figures 6.2 (buyer affected) and 6.3 (seller affected). From these, it is clear that when both parties are 'smart' (i.e. take the contract decision into account), the total utility stays non-negative, but when neither party is smart, the total utility goes negative when the effect probabilities are high. The difference between the good and bad policies is statistically significant at the $p < 0.0001$ from effect probability of 0.7 onwards in all cases and is statistically significant at least at the $p < 0.001$ level also on with probability of 0.6 and at the $p < 0.05$ level in some cases also with probability of 0.5.

In all cases, when only one of the players take the contract decision into account, one of the cases follows the case where both are smart and one follows the case where neither is smart (except that in some cases when the effect is certain, all three cases where at least one party considers the contract decision converge at the total utility of zero).[6] The case that performs better is usually the one where the seller takes the contract decision into account. However, there are two exceptions: both compensatory policy cases when the buyer is affected. There it is the buyer alone that is able to achieve non-negative total utility also at high effect probabilities. This can be explained by considering who is responsible for the seller's costs (the only source of negative total utility). In the cases where the buyer can be affected and the fee is compensatory, the buyer will always compensate for the seller's costs and therefore the seller is guaranteed at least zero profit and even if it takes the contract decision into account it will not affect its decisions because his expected utility will be automatically positive. The buyer, on the other hand, has a risk of having to pay the costs (and profits in the case of *Expectation Damages*), so when it considers the contract decision, it will be interested in adjusting its reservation price so that it will get non-negative utility on expectation. In the other cases, the buyer is not liable for the seller's costs. This is clear when only

[6]The last minute adaptations where they occur are because the agents require positive expected utility and if the utility in case of decommitment is zero but the utility in case of success is positive, the expected utility remains positive until success is no longer possible at all. When the decommitment will always happen at some point, $a = 1$, the expected utility is zero which is no longer positive and the agent will not enter any negotiations.

the seller can be affected (the buyer never pays any decommitment fees) and is also true when the buyer can be affected but decommitment is free (the *Constant 0.00* policy). Here the roles are therefore reversed, the buyer will not be interested in doing any useful adaptations whereas the seller is. This is consistent with hypothesis 9 and we can therefore accept it.

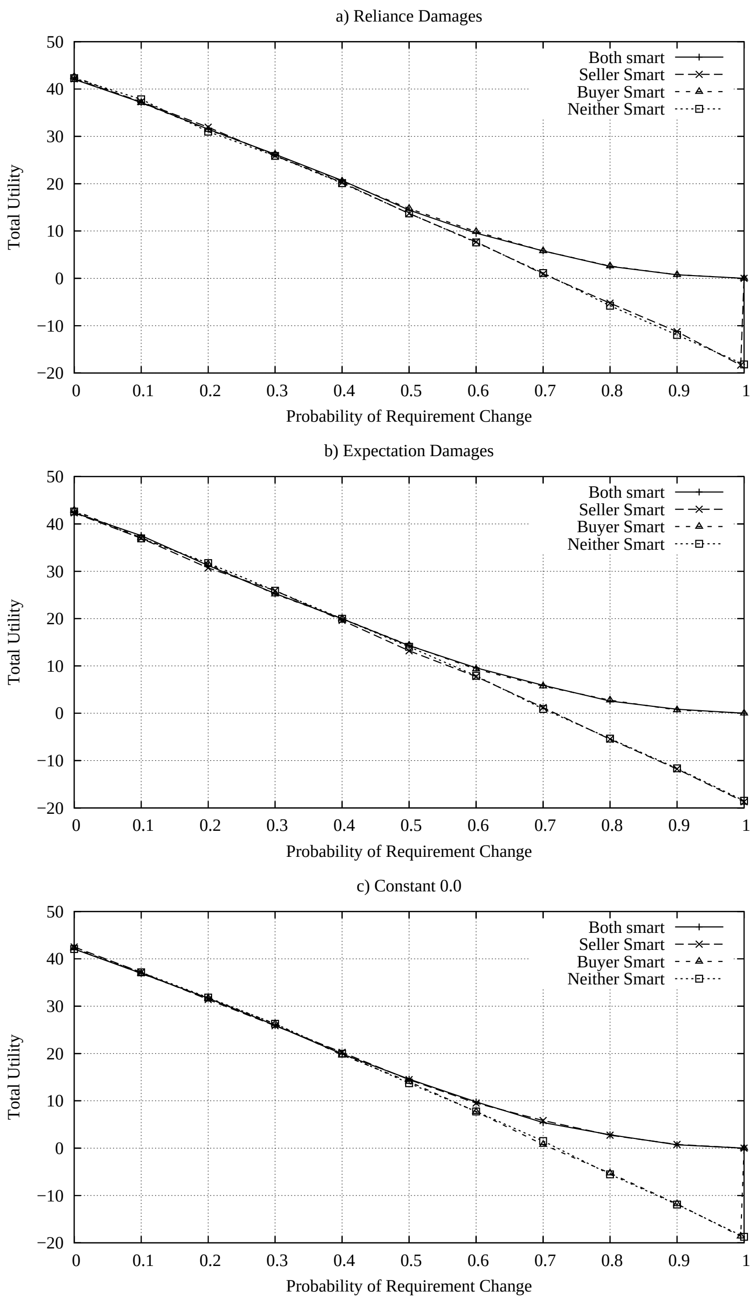

FIGURE 6.2: Contract Decision when the Buyer Can Be Affected (Hypothesis 9).

Our second hypothesis in this section considered the total utility in different cases when the decommitment policy is *Constant 1.0*. The three different situations

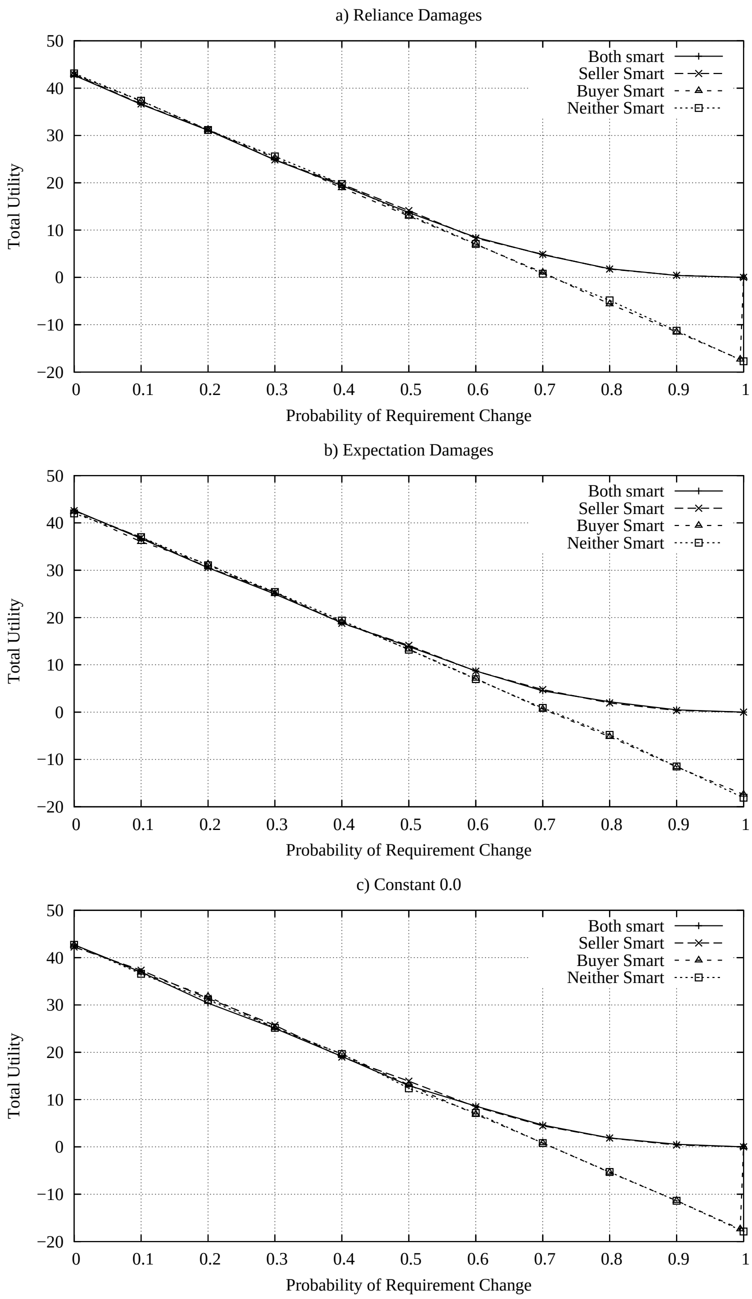

FIGURE 6.3: Contract Decision when the Seller Can Be Affected (Hypothesis 9).

are in figure 6.4. Let us start from the case where the buyer is affected (figure 6.4.a). Here, the case where the buyer takes the contract decision into account performs worse than the three other cases between the effect probabilities of 0.2 and 0.6 at the $p < 0.0001$ level except at the probability of 0.2 where it is at least at the $p < 0.01$ level. This is because the very high fee makes the buyer too cautious and the number of contracts decreases heavily. For example, when the effect probability is 0.4, the number of contracts is around 109.7 on average when the buyer alone considers the contract decision, but it is over 160 in cases where it does not and 150.9 even in the case where the seller is compensating for this attitude (both parties take the contract decision into account). The buyer's over-carefulness turns into appropriate prudence when the effect probabilities are very high and the case where only the buyer takes the contract decision into account will outperform the case where neither or only the seller will do so with effect probabilities 0.8 or greater (at the $p < 0.0001$ level). The seller's over-zealousness, on the other hand, will mean that the policy where only it will be smart, will be outperformed by the case where neither party considers the contract decision when the effect probability is 0.7 or higher. This is because the high fee gives the seller incentives to lower its reservation price so that it can get the unsuspecting buyer to enter into contracts that will very likely end in a decommitment. This will mean contracts are formed earlier and because of this there will be more of them. For example, in the case where the effect probability is 0.8, the average contract time decreases from 301.6 (in the case where neither party is smart) to 288.4 (when only the seller is smart) and the average number of contracts increases from 108.1 to 115.0 (all differences significant at the $p < 0.001$ level).

The case where the seller is affected follows the same patterns, but more strongly. This is because the seller has to consider its costs in addition to the very high fee. The case where the potential decommitter (the seller) takes the contract decision into account performs very badly compared to the other cases with intermediate effect probabilities. It is worse than all the other three with effect probabilities 0.4 and 0.5 (at the $p < 0.0001$ level) and worse than either case where both or neither take the contract decision into account between 0.2 and 0.6 (at the $p < 0.0001$ level). This is of course because the high fee and the possible costs make the seller very careful indeed. For example, in the case where the effect probability is 0.5, there are only 4.1 contracts in the case where the seller alone is smart, but 58.2 when neither is smart and 33.9 when both parties take the contract decision into account. Of course with high effect probabilities, the very careful seller strategy works better, beating the case where the contract decision is not considered by either party when the effect probability is 0.8 or above. The case where only the victim (the buyer) is smart is also much stronger here than it was when the buyer was affected. This again has to do with the seller's costs: the buyer does not consider them but the seller does. The buyer is interested in getting into contracts as soon as possible, which is often not in the seller's interest. This is evidenced by the fact that the total utility in the case where only the buyer considers the contract decision is below that of the case where neither party considers it, with

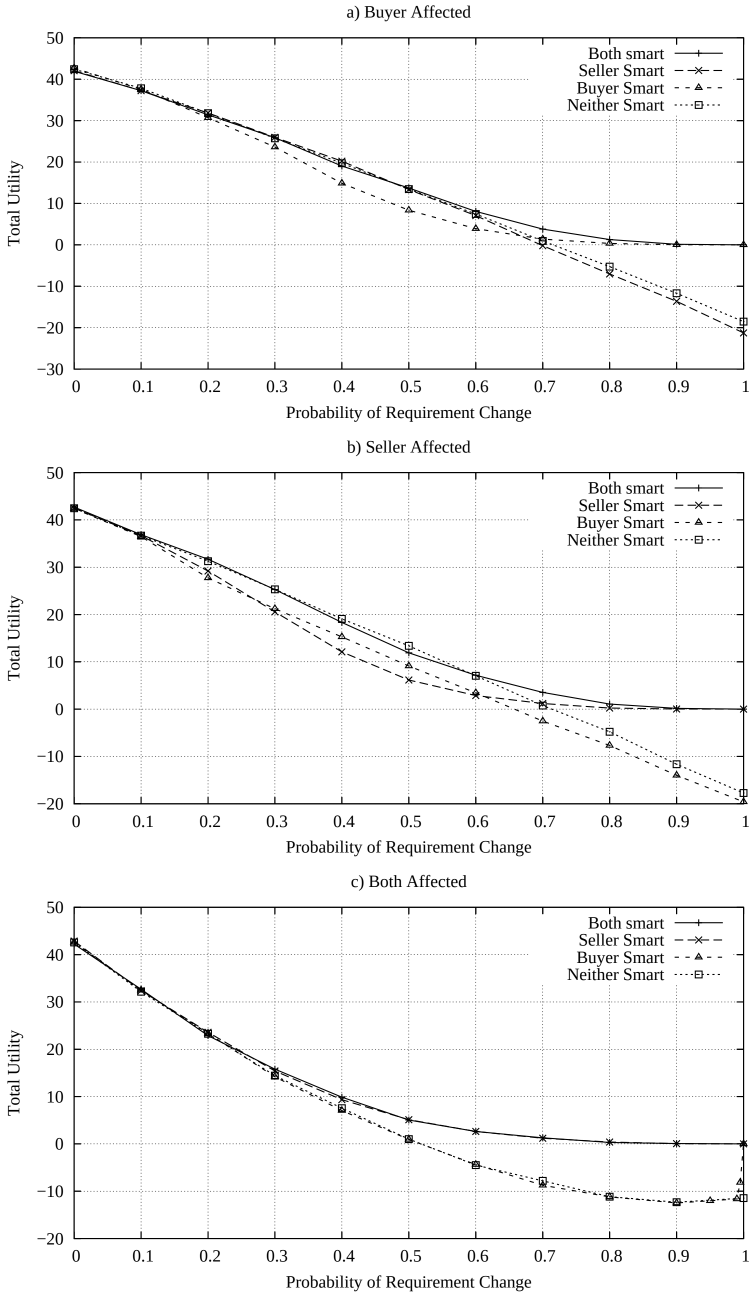

FIGURE 6.4: Contract Decision with the Constant 1.0 Decommitment Policy (Hypothesis 10).

all effect probabilities of at least 0.2 (at the $p < 0.0001$ level). For example, when the effect probability is 0.7, the average contract time is 285.9, as compared to 320.9 when neither party is smart and the number of contracts are 108.3 and 80.9 respectively. One interesting point to notice here is that the seller's aversion to contracts is so strong that even in the case where both parties are aware of the contract decision, the seller is unable to compensate and the performance in terms of total utility is worse than in cases where neither party takes the contract decision into account when the effect probability is 0.4 or 0.5 (at the $p < 0.05$ and $p < 0.001$ level respectively). All these observations are consistent with hypothesis 10 and we can therefore accept it.

We can also observe that when the risk of having to pay a high fee is more evenly distributed (as in the case where both parties can be affected with the same probability), it is easier to find a contract price that both parties can agree on and none of the peculiarities we just described occur.

The next hypothesis deals with compensatory policies when both parties can be affected. They and the *Constant 0.0* policy are shown in figure 6.5. Our claim was that the case where both parties take the contract decision into account will outperform the other cases at least in some situations. From figures 6.5.a-b we can see that this really is the case. The situations where only one of the parties considers the contract decision do better than the case where neither of them does, but there are still some cases where both of them are needed to get the best utility. Now, for *Reliance Damages* the case where both consider the contract decision beats the cases where only one of them does with effect probabilities between 0.4 and 0.8 (at the $p < 0.0001$ level between 0.4 and 0.5, at least at the $p < 0.001$ in 0.6 and at least $p < 0.05$ level in $0.7 - 0.8$). The difference here is due to the fact that both parties take the seller's costs into account only in one of the two cases and it therefore takes both of them to take them into account in all cases (as explained when we discussed hypothesis 11). Thus, there is no difference in the low effect probabilities because the risk of decommitments is too small to make a difference and in very high effect probabilities, all the best policies will negotiate very little and therefore the difference is too small to be significant.

With *Expectation Damages*, the case where both parties consider the contract decision outperforms all the other cases when the effect probability is 0.4 or above (at the $p < 0.0001$ level). This is because when only one party takes the contract decision into account, it will enter into contracts if the utility in case of success (either performance or the other party decommitting, the utility is the same) is (much) higher than the fee it has to be in case it has to decommit. The contracts will therefore be formed if the unsuspecting opponent is willing to give the smart party a contract that gives it high enough utility in case of success. This is not always possible and therefore the performance is slightly better than in the case where no one takes the contract decision into account (with effect probabilities 0.5 and above at least at the $p < 0.05$ level with probability of 0.5 and at the

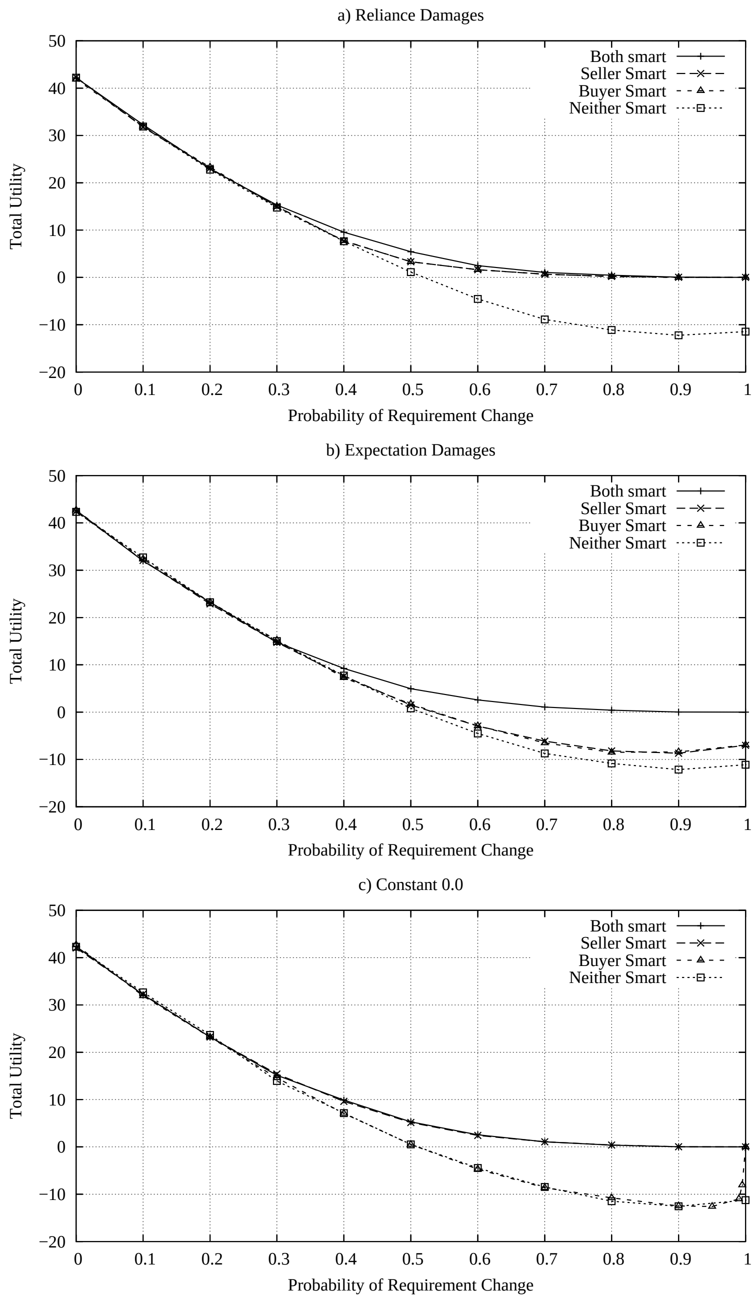

FIGURE 6.5: Contract Decision when the Both Parties Can Be Affected (Hypothesis 11).

$p < 0.0001$ level with probabilities 0.6 and above).[7] These findings are consistent with hypothesis 11 and we can therefore accept it.

Next, we are interested in the total utility under different decommitment policies. In figure 6.6 we have the four different policies in the three different settings we have (buyer affected, seller affected and both affected). The *Constant 1.0* policy

[7] Also this was discussed in more detail when hypothesis 11 was discussed in section 6.2.1.

does fare clearly worse than the other policies in the settings where only one of the parties can be affected (in the setting where only the buyer can be affected, the difference to all other policies is statistically significant at the $p < 0.0001$ level between effect probabilities $0.6-0.9$ and in the case where the seller can be affected the difference is significant at the $p < 0.0001$ level between effect probabilities $0.5 - 0.9$). With higher probabilities, all policies get very few contracts, so the difference is very small. On the other hand, when the probabilities are lower, the contract price can still be used to achieve a reasonable risk distribution. The performance of the other policies is broadly speaking the same. This is what we expected in hypothesis 12, so we can accept it.

Our last hypothesis contended that a relatively small inaccuracy, like the buyer not knowing the time the seller has to pay its cost, will decrease the total utility when the buyer can be affected and its decommitment policy is compensatory (either *Expectation Damages* or *Reliance Damages*). To this end, figure 6.7 shows the total utilities with accurate and inaccurate expected cost in the case where both parties can be affected and only the buyer takes the contract decision into account. The difference is statistically significant between the effect probabilities of $0.5-0.9$ (at the $p < 0.05$ with 0.5 and 0.9 and at the $p < 0.0001$ with $0.6-0.8$). There is similar pattern in all the other cases. This is what the hypothesis 13 suggested, so we accept it.

6.3 Summary

In this chapter, we discussed the contract decision and especially how the parties taking it into account can use the contract price as a risk-allocation tool, allowing them to find mutually acceptable contracts while ensuring non-negative expected utility for themselves and the society (contribution **C3**). In more detail, this means that both parties take possible adverse future utility changes into account by setting their reservation price so that it ensures a positive expected utility to them and through them to the society. In some situations, it is enough that only one of the parties takes the contract decision into account (when only one party is affected or when the decommitment fee is constant), but in other situations it takes both parties (compensatory policies when both can be affected) and, in some situations, one-sided consideration of the contract decision can even be harmful for the common good.

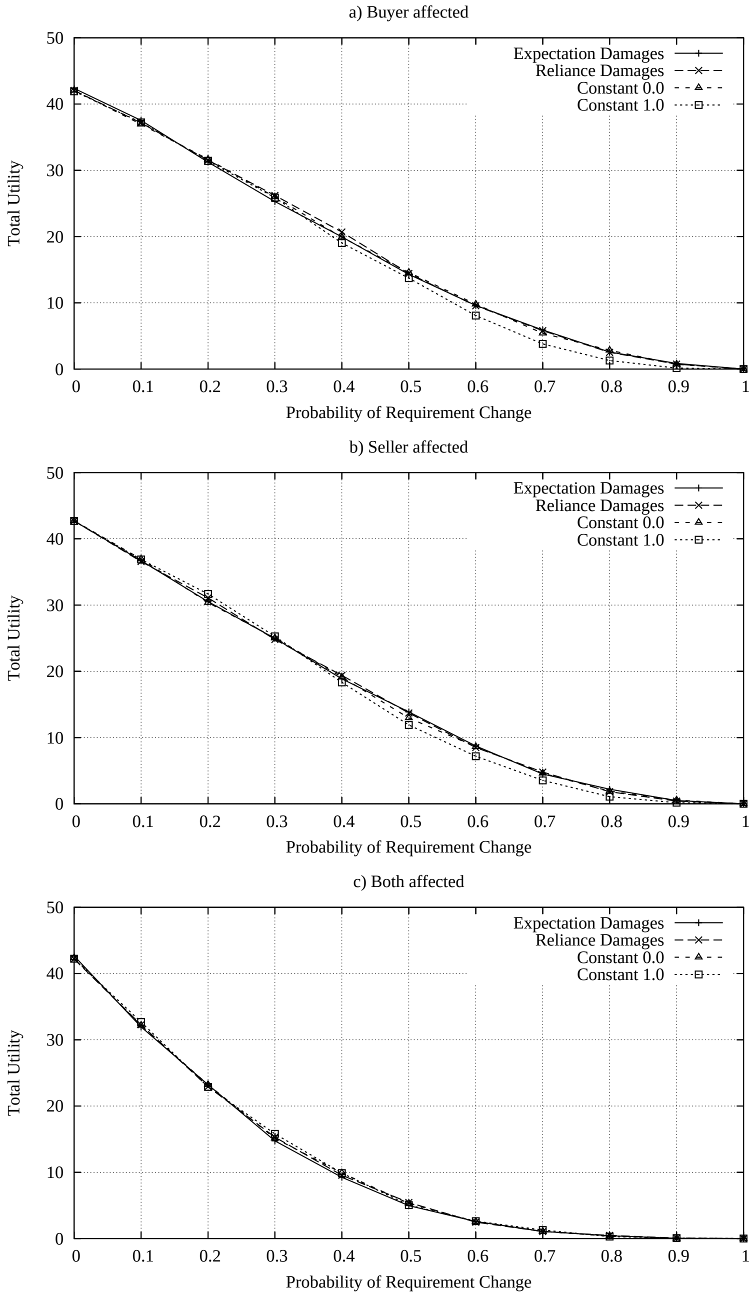

FIGURE 6.6: The performance under different decommitment policies (Hypothesis 12).

The mechanism the parties use to get these effects is to use the contract price as a means of distributing the risks between them. A high price shifts some of the risk from the buyer to the seller and a low price does the opposite. This works very well in most cases and the performance is roughly the same in all settings and under most policies. However, we found out that the *Constant 1.0* performed worse than other policies when only one of the parties could be affected. This was because we restricted our contract price to the interval $[0, 1]$ and in some situations, the parties

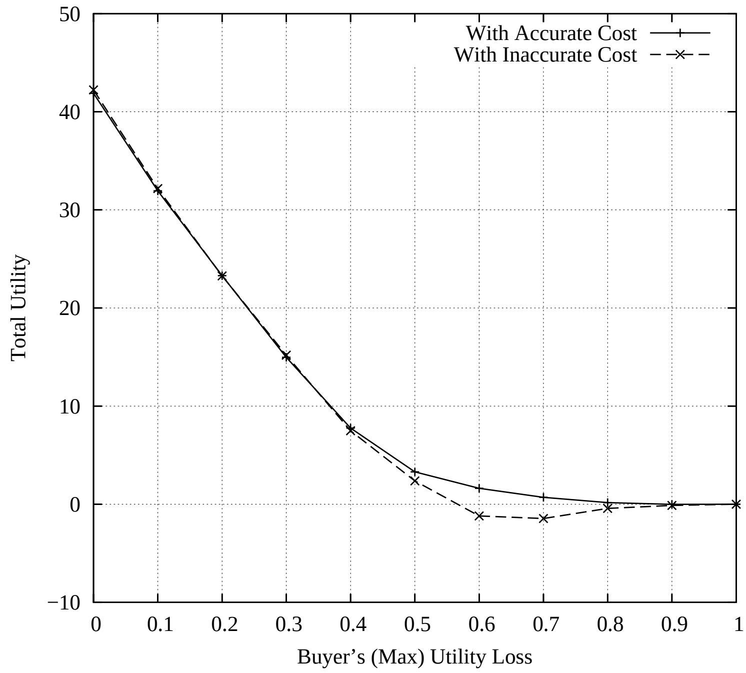

Figure 6.7: The Effect of Information Accurancy (Hypothesis 13).

are unable to do this tuning process in full. With less restrictive situations and with less extreme policies, there should be very little difference between policies.

We also showed that the quality of information plays an important role in this adaptation. Specifically, we used the buyer's information about the time the seller has to pay its costs as an example and discovered that even small inaccuracies in the information can lead to deterioration of performance, because it leads to over- or under-adaptations.

Although we investigated only a few policies, the same lessons would lend themselves to any policy. This means that all policies offering overcompensation would experience similar effects as the *Constant 1.0* did to some extent and all policies undercompensating would fare as the *Constant 0.0* policy did to some extent. So, because a policy offering an *Increasing Fee from 100% to 250% of the Price* would probably mean overcompensation for the seller, the buyer would probably be very cautious to enter into a contract, especially if the price was high. This would probably mean a lower total utility in the case where only the buyer would consider the contract decision, for example. On the other hand, from our results it is usually clear that the best total utility is achieved when both (all) parties take the contract decision into account to protect their own interests. This might have some implications on the price. For example, in the case where only the buyer could be affected and the policy offers undercompensation for the seller, the seller is likely to protect itself by asking for a higher price. Of course the policies can be arbitrarily complex: the fees could go from undercompensatory to grossly overcompensatory over time or they could be either in a single case. So the effects described might not always be self-evident from the data, but they will be there. On the other hand, we could easily have heterogeneous buyers and sellers (in terms of reliability). What it would mean is that the effects described here would be happening in each negotiation to a varying degree and they might be difficult to see from the overall data, but they will be there in each negotiation.

Next we will discuss the decision that takes place even before negotiations are even started, the selection decision.

Chapter 7

The Selection Decision

The *selection decision* takes place even before the negotiations start. It means that the party decides who to negotiate with from a number of potential opponents. For the selection decision to have an effect on common good, the opponents must, of course, be *heterogeneous*, because choosing between identical opponents obviously makes little difference. Often it is also beneficial, if the number of selectors is smaller than the number of the potential opponents, because if the numbers are roughly the same or if there are more selectors than selectees, then (almost) everybody will be selected anyway, so the selection is likely to make little difference, although as we will see this may not always be the case (if most selectees are bad and should not be selected at all and only some selection strategies take this into account).

In more detail, we investigate a setting in which the buyer chooses the sellers he wants to negotiate with and the heterogeneity of the sellers comes from two sources: they have different qualities and reliabilities. Here, as in the reliance decision (chapter 5), reliability is a probability that the seller will perform the service as agreed. The work in this chapter relates to research contribution **C4**. We will first discuss the issues raised by the selection decision and the changes we have made to the basic market model to investigate these issues (section 7.1). We will then investigate our ideas empirically (section 7.2) and conclude by summarising our results (section 7.3).

7.1 The Problem and the Modified Model

In order to investigate the selection decision, we again have made some changes to the basic market model. The biggest change is that the parties are no longer matched at random, but instead, the buyers are asked which provider they would like to negotiate with. The order in which the buyers are asked to give their choice is selected at random. When choosing negotiation partners, the obvious thing to do is to calculate the expected utility of each remaining opponent and then choose the one that has the highest expected utility. A possible contract with an opponent o has two possible outcomes: either the opponent o will perform according to the contract or o will decommit and pay the decommitment fee. The buyer utilities for these cases are as follows:

$$U_b(\text{performance}) = V(q) - p$$
$$U_b(\text{decommitment}) = f$$

And of course each and every provider has to have a reliability, ρ_s. We have the same reliability distributions as in *reliance decision*, namely *Uniform* (all values equally likely), *Normal* (average values more likely), *Exponential* (low values more likely) and *Reverse Exponential* (high values more likely). The reliability will be drawn at random from the distribution separately for each provider s (reliabilities are independent from each other). When the contract has been entered into, the parties exit the market and a random variable (from the distribution $Uniform(0,1)$) is drawn and if this value is greater than the seller's reliability, the seller will have to decommit at some point and this point is selected at random from $Uniform(t_{contract}+1, t_{delivery}-1)$. This means that unlike in the basic model, the probability of the effect is different for each seller ($a_s = \rho_s$ varies). This is done to enhance the meaning of opponent selection. On the other hand, here any effect will always lead to decommitment. We ensured this by setting the effect $l_s = 5.0$. This change effectively removes the performance decision from the model.[1]

Since each seller's reliability, ρ_s, gives the probability that the provider is going to perform, the expected utility for the buyer b to negotiate with the seller s is:

$$\begin{aligned} U_b &= \rho_s U_b(\text{performance}) + (1-\rho_s) U_b(\text{decommitment}) \\ &= \rho_s (V(q) - p) + (1-\rho_s) f \end{aligned}$$

[1] If in some cases, the parties would not decommit in case of effect, they would make a performance decision (whether or not to perform). When we always decommit in case of an adverse effect, there is no performance decision to speak of.

Of course the problem in terms of this calculation is that the negotiation outcome (negotiation price p) is uncertain before the negotiation has started and since we assume that here, as in the basic model, the parties' β parameters are selected at random (see section 3.2) it can vary significantly. However, it is clear that the contract price will be in the range $[C(q), V(q)]$ because nothing else would be accepted by both parties at the same time.[2] On the other hand, since the β will be selected at random and independently for both parties, it seems that the most reasonable assumption for the average or expected outcome is the average of these two figures, so the expected price is:

$$Ep = \frac{V(q) + C(q)}{2}$$

We therefore get:

$$\begin{aligned} U_b &= \rho_s(V(q) - Ep) + (1 - \rho_s)Ef \\ &= \rho_s(V(q) - \tfrac{V(q)+C(q)}{2}) + (1 - \rho_s)Ef \\ &= \rho_s \tfrac{V(q)-C(q)}{2} + (1 - \rho_s)Ef, \end{aligned}$$

where Ef is the expected fee, which of course depends on the decommitment policy in use. Here, we use very simplistic estimates (for convenience). For example, in all types of increasing fee policies, we use an average of fee in the next round (assuming the contract was formed at the time of selection) and the full fee just before the $t_{delivery}$.[3] Therefore we do not use any information about *when* the possible loss is going to take place and the actual fee can be larger or smaller depending on when the decommitment takes place and what the actual contract price is. What matters at this stage is only the expectation. We will use the 106 of the 107 policies described in section 3.3. We do not use the *Reliance Damages* policy here, because given that the buyer does not have any costs, it will be equal to the *Constant 0.0* policy and so, we use that instead.

There is a set of cases where no price that both parties would agree on can be found in the basic model. That occurs when the quality of the service is lower than the buyer's minimum quality, q_b^{min}. This means that $V_b(q) = 0$, but $C_s(q) > 0$. Therefore negotiating with opponents that do not provide at least the minimum

[2] Since the negotiation after the selection uses the basic model, the buyer will not accept any offers in the negotiation that would be higher than the value of the service, $V(q)$, and on the other hand, the seller will not accept offers that are lower than their costs, $C(q)$. This means that the price, an average of the two, must be between $C(q)$ and $V(q)$.

[3] For the *Constant* policy variations, the expected fee is obvious. In the *Expectation Damages* policy, the estimated fee will be $\frac{V(q)-C(q)}{2}$.

utility for the buyer is a waste of resources and all buyers ignore these sellers.[4] Of course the minimum quality varies from one buyer to the next and the provider that is ignored by one buyer might not be by another. However, it is clear that the providers with very low quality will have trouble finding buyers to negotiate with (no matter what the decommitment policy is). The decommitment policy does not affect this restriction.

Another restriction we have in the selection is that the expected utility has to be positive for a buyer to consider an opponent at all. This is always the case when $V(q) \geq C(q)$ and the fee f is non-negative, so with any policy that has non-negative fees, this restriction never plays any role. In the basic model (and the 106 policies we got from it), the fees f are always non-negative, but here we also allow negative fees. In this context, a negative fee means that instead of a seller paying the buyer in case of decommitment, it will be the buyer paying the seller (even if the decision to decommit was seller's). The reason for allowing negative fees is that the expected utility above is only the expected utility of the buyer and it does not take into account the seller's or the society's welfare. From the seller point of view, only a performance will be able to bring positive utility to it. In the case of decommitment, it will have to pay the fee and possibly also the preparation costs it has paid at the time of decommitment. For the society, the fees do not have an effect on the total welfare: No matter what the decommitment fee is, it will only be moving from one party to another and it will not generate any new welfare. However, decommitment may lead to losses to the society, since the seller might have (with probability $D(t)$) paid its cost (c_s). So the total utilities are:

$$\begin{aligned} U_{b+s}(\text{performance}) \quad &= U_b(\text{performance}) + U_s(\text{performance}) \\ &= (V(q) - \tfrac{V(q)+C(q)}{2}) + (\tfrac{V(q)+C(q)}{2} - C(q)) \\ &= V(q) - C(q) \end{aligned}$$

$$\begin{aligned} U_{b+s}(\text{decommitment at turn } t) \quad &= U_b(\text{decommitment}) + U_s(\text{decommitment}) \\ &= f + (-f - D(t)c_s) = -D(t)c_s \end{aligned}$$

[4] Of course if the buyers would consider contract decision during the negotiation, it might be possible to find a contract even in these cases, since the buyer could offer a positive price for the service in hope of non-performance and a (hugely positive) decommitment fee. They would then choose very unreliable providers to negotiate with. However, if the sellers take the contract decision into account too, the very unreliable providers might not show up in the marketplace at all or would require a much higher price.

And the expected total utility is:

$$U_{b+s} = \rho_s(V(q) - C(q)) + (1 - \rho_s)(-D(t)c_s)$$

From society's point of view, it would therefore be useful, if the buyers selected more reliable providers. However, this does not necessarily mean the most reliable providers, because also the utility of the contract (both when performed and when decommitted from) matter. It may well be that a buyer (and indeed the society) prefer some less-than-perfectly reliable provider that can offer excellent utility in the case of performance to a perfectly reliable seller that can only provide a meagre utility. And although higher quality providers usually provide better utility when they perform, they might also suffer bigger losses if they fail. It is therefore the combination of these factors that should be optimised.

When the buyer decides which seller to negotiate with, it will only consider its own utility and choose an opponent that will provide the best expected utility for it. The main difference between the buyer's and total utility is that in the case of decommitment, the buyer gets a decommitment fee of f, but the society faces a loss ($D(t)c_s$), because the seller has to provide the fee in question plus pay its possible costs. As in the case of the performance decision, the optimal decision from the society's point of view would be the one that internalises the costs of the opponent to the buyer's decision-making. Here, it would mean paying for the seller's possible costs. However, there is another difference between the buyer's and the society's utility. In the case of performance, the buyer expects a utility of $\frac{V_b(q)-C_s(q)}{2}$. That is, a half of the society's benefit (the other half was assumed to belong to the opponent and the buyer does not consider that). So, we get the following two decommitment policies:

- *Reverse Reliance Damages* policy, where the buyer would, instead of getting a fee, have to compensate the seller for its costs, although it is the seller that decommits. In other words, the fee would be $-D(t)c_s$.
- *Enhanced Reverse Reliance Damages* policy, where the buyer would, instead of getting a fee, have to compensate the seller for its costs, although it is the seller that decommits. In addition, the 'seller's' utility in the case of success is taken into account. The fee is $-D(t)c_s + \frac{\rho}{1-\rho}\frac{V(q)-C(q)}{2}$.

The latter makes the buyer's utility converge to that of the society's. Therefore each buyer will make an optimal choice for the society on its own turn. However, there are two main concerns with this policy. The first is that although it ensures that each player at its turn will choose the best opponent for itself from the society's point of view, the order in which the buyers will make this choice is still selected at random and it may well lead to sub-optimal matchings. This is, for example, because the same seller may be the best opponent for many buyers and only one of them (selected at random) will be able to negotiate with it. On the other hand, the restrictions ($q > q_b^{min}$ and $U_b > 0$) may mean that sometimes all the sellers that could be chosen are chosen. To see why this is, consider a case where there are three buyers and three sellers. The sellers have qualities 0.2, 0.3 and 0.4 and the buyers have minimum qualities of 0.15, 0.25 and 0.35. Now, in the optimal ordering the buyers would be given a choice in the reverse order: 0.35, 0.25 and 0.15. This would mean that the first one would choose the seller with $q_s = 0.4$, the next one the seller with $q_s = 0.3$ and the last one the seller with $q_s = 0.2$. Everybody would be able to find a negotiating partner and, given that $V(q) > C(q)$ is guaranteed to hold, a contract. However, if the order was reversed, then the buyer with $q_b^{min} = 0.15$ would choose first and he could take any of the three opponents because they all have qualities above his minimum, but if he chooses something other than the first one then either buyer 2 or 3 will be left without a suitable negotiation partner.

The only way around this problem would be to choose the order in some more sophisticated way or let the marketplace optimise the match-up. However, most of these problems apply to all policies and over many repetitions. Given this, the different possible orders average out leaving the parties with something between the best and worst case on expectation. Therefore, on expectation, the *(Enhanced) Reverse Reliance Damages* policies should, however, do better than any policy with non-negative fees, especially when there are plenty of providers to choose from and many of the providers are relatively unreliable. Of course also getting those actual costs might be difficult in any non-trivial setting (as we discussed in the context of the performance decision). Another problem with these policies is more fundamental: They would mean that the victim would have to compensate the decommitter's costs and that might be quite problematic, not least because it might give the wrong incentives to the seller. We will discuss this problem further in the summary of our results (section 7.3).

When a policy allows negative fees (the victim paying for the decommitter for his decommitment), the expected utility for a given opponent might be negative and since only opponents with positive expected utility are negotiated with, this may mean that the policies that offer negative fees may not select all opponents even if the number of sellers is bigger than that of buyers. When negative fees make the buyer's utility reflect the total utility, as in the case of *Enhanced Reverse Reliance Damages* policy, the buyers will choose only the sellers that will provide them and the society with positive expected utility. However, in other cases, the negative utility may mean that some potentially useful opponents are not negotiated with, which may be problematic. To investigate this, we will also introduce the following policies:

- *Reverse Expectation Damages* policy, where the buyer would, instead of getting a fee, have to compensate the seller for its loss (profit + costs), although it is the seller that decommits.
- *Constant* policy, where the buyer would, instead of getting a fee, pay the seller (the decommitter) a constant fee fee. We use fees -0.25, -0.50, -0.75 and -1.00.

These policies will make decommitment a very unattractive option for the seller and therefore increase the importance of reliance to the buyer. This means that the buyers prefer reliable opponents even if the utility they provide in case of success is very modest and if no reliable providers are available, these policies encourage the buyers not to negotiate at all. This means that they will ignore some negotiations that could lead to positive utility for the society and therefore these policies are likely to do worse than the *(Enhanced) Reverse Reliance Damages* policies.

A policy that would be easier to implement is one that would make decommitment free for the seller (the *Constant 0.0* policy). This is because any positive fee would mean that the buyer would be less interested in negotiating with more reliable providers. In one special case, the *Expectation Damages* policy, the buyer would be indifferent between performances and non-performances (he gets the same utility anyway) and therefore he would select the opponents that can provide the best expected utility in case of success, no matter what their reliability is. In addition, if the fee is actually bigger than the buyer's utility for the performance, the buyer actually prefers non-performance and chooses more unreliable providers. Both of these strategies are problematic from the society's point of view, because as explained, only a performance can improve the total utility (i.e. the common good). On the other hand, the zero decommitment fee of course means that the buyer will not take the costs of the seller into account when choosing the sellers. This is problematic because the providers that provide the best value (best quality) also usually have the highest costs. But this may not be critical, especially if the reliability of those chosen sellers is reasonably high (and, therefore, so is the probability that they would have to pay their costs without getting anything in return). The fact that the buyer only considers his half of the common good in case of success is less problematic because a half of any number is bigger than a half of any other number only if the latter number is bigger and its relation to the utility in the non-performance case (=zero) will always be the same.

Now, the relative sizes of buyer and seller populations are likely to have a strong impact on the total utility. When the number of providers is significantly higher than that of the consumers, the providers the consumers choose has a bigger impact on the total utility, because many providers will not be selected at all and therefore the ones you do choose make an impact.[5] To investigate the effect of population size differences, we will use the *Medium* population size for the buyers and the full range of population sizes (from *Tiny* to *Huge*) for the sellers. The *Medium* population size for the buyers was chosen so that we can observe how the total utility behaves in all three situations: the number of buyers is smaller, the same or larger than the number of sellers. In addition, the *Medium* size also gives us a large enough number of contracts to clearly see the differences between different situations.

7.2 Empirical Evaluation

Now, we have described the theory, but we still need to show that the theory works in practise. We will first derive some hypotheses from the discussion above (section 7.2.1). Then we will discuss some details of how our experiments are run (section 7.2.2). And then we will discuss our results (section 7.2.3).

7.2.1 Hypotheses

From our discussion above, it is clear that the *Enhanced Reverse Reliance Cost* and *Reverse Reliance Cost* policies should be very useful in terms of total utility. The *Enhanced* version makes the buyers choose the best seller for the common good and the basic version is not far from that since it takes the losses the decommitment

[5] If the populations are roughly the same size or if there are more sellers than buyers, (almost) all providers will usually be chosen and so, the selection itself would have a limited effect (especially since the order the consumers are allowed to choose is selected at random, so if the populations would be of the same size, the difference to random matching would be relatively small).

causes to the seller (and society) into account. The only difference between the two is the fact that the seller's utility is also considered in the case of successful performance in the enhanced version. This difference was $\rho_s \frac{V_b(q)-C_s(q)}{2}$ and this is clearly increasing in reliability ρ_s. However, this 'correction term' has very limited meaning if the ρ_s is very high, because then the meaning of the non-performance utility is very small and it matters little if the expected utility in the success for case is $V_b(q) - C_S(q)$ or $\frac{V_b(q)-C_s(q)}{2}$. The difference only matters when the non-performance is clearly possible. On the other hand, if ρ_s is very low, then the difference between *Enhanced Reverse Reliance Damages* and *Reverse Reliance Damages* policies is very small. This seems to indicate that these policies choose different opponents only if ρ_s is somewhere in the middle or around 0.5. This means that there is a difference only if the consumers have to or are able to (depending on the circumstances) choose (also) providers with intermediate reliabilities. This occurs when either the average reliability of the sellers is low (*Exponential* distribution), when the number of more reliable providers is limited compared to the number of sellers (*Normal* distribution) or when the reliability varies and the seller population is not that much larger than the buyer population (*Uniform* distribution). In other words:

> **Hypothesis 14.** The *Enhanced Reverse Reliance Damages* is no worse than *Reverse Reliance Damages* and in cases, where a significant number of providers with moderate reliability (around 0.5) are chosen, it will be better.

In a similar fashion, we can reason about the *(Enhanced) Reverse Reliance Cost* policy's performance against *Constant 0.0* and all other policies with non-negative decommitment fees. These policies do not take the costs of the seller into account. This means that when the probability of this cost increases (the reliability decreases), they will perform worse. We therefore contend:

> **Hypothesis 15.** The *(Enhanced) Reverse Reliance Damages* is no worse than any other policy and in cases, where also relatively unreliable sellers have to be selected, it will be better.

Now, when it comes to the *Constant 0.0* policy and the policies with positive decommitment fees, none of them take the sellers' costs into account, so the only source of difference (if any) comes from the set of providers they choose and especially their reliabilities. To see see why this is, consider a population where all sellers have the same reliability. In this case, the expected utility of all sellers would have a term $(1-\rho_s)f$ and given that $\rho_s = \rho$ for all sellers, this would mean that this term is constant given a decommitment policy, assuming that the fee does not depend on the price. In such a case, the opponents would always be selected based on the performance utility, $\frac{1}{2}(V_b(q) - C_s(q))$, alone. And if the fee depends on the price, one should notice that the higher the quality, the higher the expected price and therefore the expected fee. Since the fee will be positive from the buyer's point of view, that makes the high quality providers even more interesting for the buyer.[6]

Therefore the main source of difference between these policies is mainly in the reliability of the selected opponents. And the higher the fee awarded in case of non-performance, the less worried the buyer is about non-performance. Up until the fee reaches the success utility, the buyer will be more and more interested in choosing opponents that produce good results in case of success. When $f = U_b(q, p|\text{success})$ (*Expectation Damages* policy), only the success utility of the opponent matters, because that is the utility also in case of non-performance. When the fee exceeds the utility in case of performance, the buyer starts to prefer less reliable providers more and more. Now this latter behaviour is especially devastating to the common good, because the probability of non-performance increases and non-performance cannot improve common good but instead it often decreases it (due to the seller's costs).

This means that the *Constant 0.0* policy should be the best among the policies that have non-negative fees and it should clearly beat all policies that in a significant number of cases choose less reliable providers (the fee is higher than the success utility for the buyer). This includes all policies that have fees that are significantly higher than zero. However, when the fees are close to zero, the difference to the *Constant 0.0* policy is going to be relatively small. Of all the policies that we will try here, we consider *Increasing 0.00-0.25*, *Increasing Contract Time 0.00-0.25*, *Increasing Decommitment Time 0.00-0.25*, *Increasing Price 0-50%* and *Increasing Contract Time Price 0-50%* and *Increasing Decommitment Time Price 0-50%* to have relatively low fees. The opponents chosen are only slightly (if at all) different. This means that *Constant 0.0* policy may not always be able to beat such policies with low fees in a statistically significant way. This is because in those cases, the differences in fees and expected utilities (and therefore probably also in the opponents the agent ends up choosing) between *Constant 0.00* and other policies

[6] On the other hand, with the quality also the expected cost is likely to rise, so that may not be optimal from the society's point of view.

are relatively small. However, this depends on the reliability distribution. If the reliabilities can be widely different or there is a significant number of low reliability providers, some less reliable providers may have quite high expected utilities if they provide good utility in case of success and some positive utility (a positive fee) also in case of decommitment. This follows from the fact that the lower reliance is compensated for to some degree by the positive fee in the case of non-performance and therefore choosing less reliable providers is less risky. The differences are larger when the reliabilities overall are low (like in *Exponential* distribution) or there is a significant proportion of low reliabilities (as in the *Uniform* distribution). The effect is less likely when the majority of reliabilities are broadly similar (like in the *Normal* distribution) or when the reliabilities tend to be high (in the *Reverse Exponential* distribution). When the fees are significant (clearly larger than zero), the difference should be clear in all of our settings. Thus, we contend:

> **Hypothesis 16.** The *Constant 0.00* policy is no worse than any policy that uses non-negative decommitment fees and it will be able to beat all policies that use high decommitment fees in all our settings and it will also beat policies with relatively low decommitment fees in some settings.

Our final hypothesis deals with other policies with negative fees. These are *Reverse Expectation Damages* and the *Constant* policies with negative fees. These policies will of course be no better than the *Enhanced Reverse Reliance Damages* policy and will be beaten by the optimal policy in many situations. However, we will be more interested how these policies fare against the *Constant 0.0* decommitment policy. All these policies are likely to over-compensate for the seller's losses and especially his costs. Of course the costs may be over 0.25 in which case *Constant -0.25* will under-compensate for the costs, but often even that policy will overcompensate. This means that the buyer is very worried about possible decommitments and will try to choose more reliable providers and if there are not reliable enough providers in the market, it will prefer not to negotiate at all. This will mean that these policies will negotiate less than other policies and they will also pass on the opponents that could bring positive utility for the society. On the other hand, since these policies usually overcompensate, the sellers usually have positive utility in the end no matter what happens, unlike under the *Constant 0.0* policy where the sellers may often be liable for their wasted preparation costs. When the *Constant 0.0* policy chooses a large number of relatively low reliability providers, many of them may be detrimental to the common good and therefore very careful policies with negative fees may outperform the *Constant 0.0* policy in terms of total utility. However, when sellers with more intermediate reliability are chosen, the best policies with negative fees are likely to be too careful and therefore fare worse than *Constant 0.0*. When reliabilities are high all around, there is likely to be little difference between the two.[7] Thus, we contend:

> **Hypothesis 17.** The *Constant 0.00* policy will do well against the *Reverse Expectation Damages* and *Constant* policies with negative fees when many sellers with moderate or high (not very high) reliability are chosen, but less well when very bad providers have to be chosen. When reliabilities are very high, the difference (if any) is likely to be small (decommitments are rare).

7.2.2 Experimental Setup

In all our experiments we had a *Medium*-sized buyer population. We then ran the experiments in four different reliability distribution settings (Uniform, Normal, Exponential and Reverse Exponential) and with seven different seller population sizes (tiny, very small, small, medium, large, very large, huge). That is 28 settings in total. In each setting, we ran the market using *(Enhanced) Reverse Reliance Damages* and all 107 basic decommitment policies. Each run was repeated 100 times and averages of total utility were calculated. A statistical analysis was then conducted usually using one-sided t-test to compare the averages. However, in the testing of the final hypothesis, we use two-sided tests, because it would be difficult from our theory to know exactly where the changes described occur.

7.2.3 Results

The performance of different policies with different sizes of seller populations and different seller reliability distributions is shown in figure 7.1. It is clear in all four reliability distribution cases that the *Enhanced Reverse Reliance Damages* policy is practically always on the top and the *Best of the Rest*, the best of all policies that do not have low fees (all 101 of them) is the worst when the seller population

[7] All of these claims require that the negative fees are reasonably large and that is why -0.25 is the smallest negative fee in the *Constant* policies. If the negative fees could be very small, the situation would be analogical to the situation where there is small positive fees: the differences would be smaller.

size exceeds the buyer population size and the selection starts to have an effect on the performance. It is also easy to see that in all cases when the seller population is *Medium* or of smaller size, the performance of all policies except the *(Enhanced) Reverse Reliance Damages* is broadly similar.

The reason why the *(Enhanced) Reverse Reliance Damages* policies are better than the rest, even when the seller population is smaller than the buyer population, is that they only start negotiations that will lead to positive expected utility for the society. This is because they also take into consideration the seller's costs. The other policies do not and therefore they may start negotiations that would not be in the society's interest, although they are in the buyer's. The difference is especially clear when the reliabilities are low (reliabilities *Exponential*ly distributed, see figure 7.1.c). Here, due to the low reliabilities, many contracts are detrimental to the common good and the total utility in most policies actually decreases when the population size and the number of contracts formed increases. This trend only changes when the opponent selection starts to kick in (the seller population is larger than the buyer population) and only more reliable providers are selected.

As for our hypotheses, we claimed that the *Enhanced Reverse Reliance Damages* policy will be able to beat all the other policies (including the *Reverse Reliance Damages* policy) when a large number of moderately reliable providers are among the selected. There is a statistically significant difference in performance of these policies with:

- *Uniform* reliability distribution when the seller population is *Medium* (at the $p < 0.01$ level),
- *Normal* reliability distribution when the seller population is *Medium*, *Large* or *Very Large* (all at the $p < 0.01$ level) and
- *Exponential* reliability distribution when the seller population is *Large* or *Very Large* (all at the $p < 0.05$ and $p < 0.001$ level respectively).

In all other cases, their performance is statistically inseparable. With the *Uniform* distribution, there is a large number of providers with moderate reliances chosen when pretty much all providers are chosen (Medium). When the seller population size increases, both policies discussed here find more and more very reliable providers and the difference vanishes. The same effect can be seen in other cases. In the *Normal* distribution cases, most of the providers will be in the reliability range that makes the difference and the number of very reliable providers grows

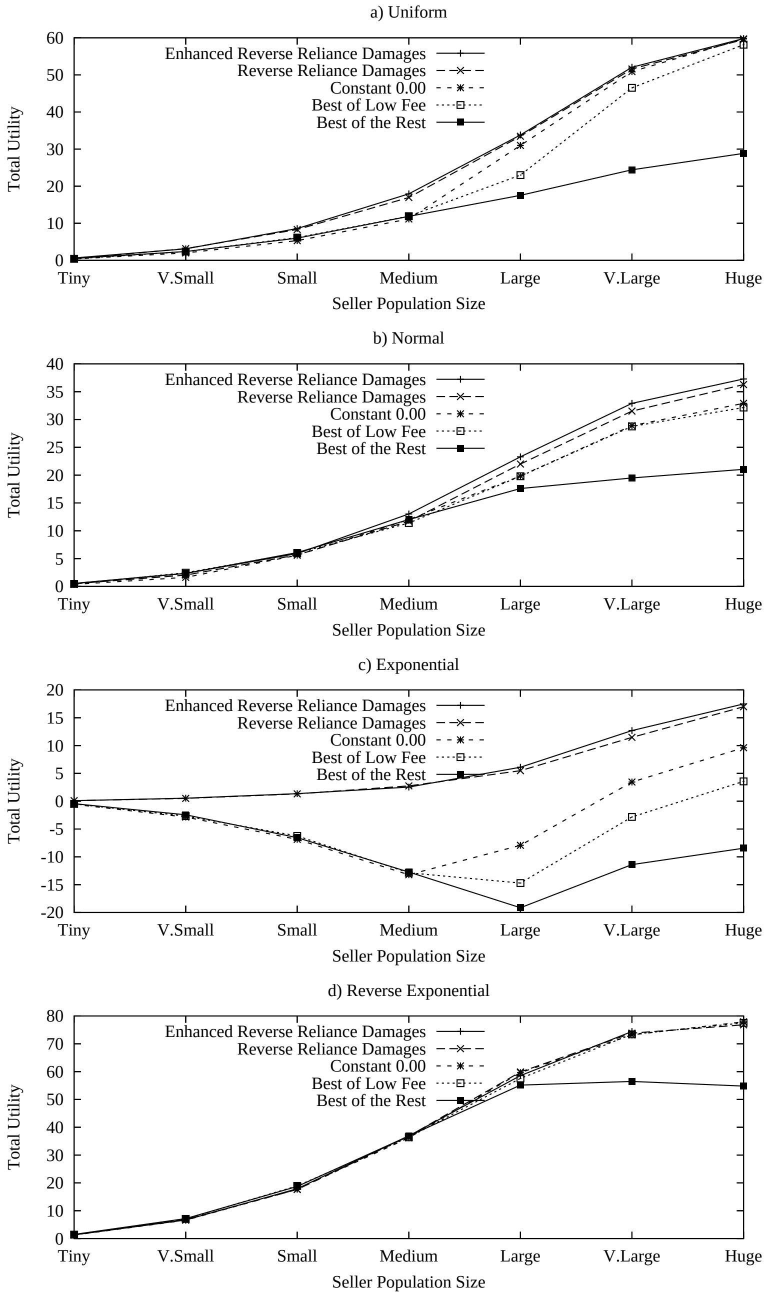

FIGURE 7.1: Selection Decision (Hypothesis 14).

slowly when the number of sellers increases. Therefore, the difference is clear up until *Very Large* population size.[8] In the *Exponential* distribution case, there are too many very unreliable providers chosen in the *Medium* population size case. But when the population size increases further, there are more moderately reliable providers in the market, they are chosen and the difference is clear until the number of relatively reliable providers increases in the *Huge* population to such

[8]And actually the p-value also in case of Huge population is 0.0504, so making it very close to being significant.

an extent that the difference vanishes again. In the *Reverse Exponential* reliability distribution cases, there are simply too many very reliable providers to show any difference. These findings are consistent with hypothesis 14 and we therefore accept it.

The next hypothesis argues that the difference between the *Reverse Reliance Damages* and *Constant 0.0* policy is at its largest when there is a large number of relatively unreliable providers in the seller population and the buyers are forced to choose them. There is a statistically significant difference between the *Enhanced Reverse Reliance Damages* and *Constant 0.0* policies in the following cases:

- *Uniform* reliability distribution, all seller population sizes except *Huge* (*Tiny* at the $p < 0.01$ and *Very Large* at the $p < 0.05$ level, otherwise at the $p < 0.0001$ level),
- *Normal* reliability distribution, all seller population sizes (from *Large* upwards at the $p < 0.001$ level, *Very Small* at the $p < 0.001$ level and in other cases at the $p < 0.05$ level),
- *Exponential* reliability distribution, all seller population sizes (at the $p < 0.0001$ level).

There is no difference in the other cases. In cases where the seller population is at most *Medium*, the differences are largely explained by the fact that the *Enhanced Reverse Reliance Damages* policy only chooses opponents that will produce positive expected utility for the society. The *Constant 0.0* does not take the seller's costs into account, so it will be less careful and that will decrease its total utility. In the *Uniform* distribution, the *Constant 0.0* policy will choose some relatively unreliable providers until there are so many reliable providers to choose from (the seller population is *Huge*) that it manages to choose enough very reliable ones that there is no difference. In the *Normal* and *Exponential* distributions, there is never enough of these very reliable sellers, so the difference remains in all cases. In contrast, in the *Negative Exponential* case, there is always so many very reliable providers that there is no difference between the policies. The difference between the *Constant 0.0* and *Enhanced Reverse Reliance Damages* policy is much larger than between the *Reverse Reliance Damages* and *Enhanced Reverse Reliance Damages*, so the differences are also clearer in the other cases and the differences are especially clear in the *Exponential* distribution case where the reliabilities are in general very low. This is what hypothesis 15 claimed, so we are able to accept it.

From figures 7.1.a-d, it is clear that the *Constant 0.0* policy beats all policies with high fees easily when the opponent selection affects the outcome (when there are more sellers than buyers). This difference is at the $p < 0.0001$ level in all cases except in the *Normal* distribution and *Large* seller population size where it is at the $p < 0.001$ level. When it comes to the policies with low fees, the differences are much smaller. The cases where there is a difference are as follows:

- *Uniform* reliability distribution, from seller population size of *Large* onwards (at the $p < 0.01$ level in *Huge* population case and at the $p < 0.0001$ in the two other cases),
- *Exponential* reliability distribution, from seller population size of *Large* onwards (all at the $p < 0.0001$),
- *Reverse Exponential* reliability distribution, the seller population size of *Large* (at the $p < 0.001$).

In the *Normal* distribution there was no difference and in cases where the seller population was smaller, all policies with non-negative fees were producing broadly similar results. In *Uniform*, there is a wide range of different reliabilities among the sellers and some sellers will be relatively unreliable but offer a large utility in case of performance. This tempts the buyers to choose them more when the policy gives them positive utility even in case of non-performance. In the *Exponential* case, there is a large number of unreliable providers and the buyers are less worried about choosing them if they get some positive utility also in case of non-performance. On the other hand, with the policy *Constant 0.0* there is no benefit for the buyers in non-performance, so they will tend to choose more reliable providers.[9] In the *Reverse Exponential* case, there is still not so many providers that have both a very good quality (value) and are very reliable when the seller population is *Large* and therefore the policies with positive fees tend to choose some providers with lower reliability. However, when the population size increases, there is enough high quality and high reliability providers to choose from such that the difference will not be statistically significant. The reason why there is no difference in the *Normal* distribution is that in this case the vast majority of providers have relatively similar reliability (around 0.5) and therefore the policies with low positive fees are less tempted to choose very unreliable providers.

[9] Of course, the *Constant 0.0* policy will not choose the *most* reliable providers because $\rho_s \frac{V_b(q)-C_s(q)}{2}$ still has two factors and if the $\frac{V_b(q)-C_s(q)}{2}$ is large enough a slightly lower ρ_s will do. But with a positive fee also a term $(1-\rho_s)f$ will compensate for the lower ρ_s and therefore the buyers will under these policies choose even more unreliable providers.

The performance of policies with negative fees are shown in figure 7.2. As expected, the clearest cut case is the *Normal* distribution, where there are very few very reliable providers and many others mediocre reliabilities. In these circumstances, the negative policies that try to avoid decommitments by selecting very reliable providers and if they fail that not negotiate at all, do not perform well and they are clearly worse than the best policies. With very high negative fees (*Constant -0.5* or *Constant -1.0*), there is barely any negotiations at all. The *Constant 0.0* policy is able to beat *Reverse Expectation Damages* and all *Constant* policies with negative fees in all cases except *Tiny* (with *Very Small* population at the $p < 0.001$ level and in other cases at the $p < 0.0001$ level[10]). Another clear-cut case is the *Exponential* distribution, where there are many unreliable providers and many of the providers are actually counterproductive to the common good. All discussed policies with negative fees get non-negative values in all situations, though, because they only negotiate with opponents that can give them positive expected utility and their fees are likely to cover the seller's costs in any case. This means that the expected utility is non-negative for the society too. However, the buyers are negotiating very little especially when the negative fees are very high. This careful attitude pays in this setting and the best policies are able to beat the *Constant 0.0* policy in all seller population cases except *Huge*. In the *Huge* seller population, the *Constant 0.0* policy manages to select so many good providers that it is no longer worse, but it will not be able to beat the *Reverse Expectation Damages* here in a statistically significant way ($p = 0.0694$).[11]

The remaining two cases are less clear, although the story is similar. In the *Uniform* distribution, there is a reasonable number of providers with very good reliabilities and therefore most of the policies with negative fees will be able to perform well (the ones with very high fees obviously will perform less well). Here, the *Constant 0.0* policy overtakes the negative fee policies already when the seller population is *Large*. This is because the opponents it chooses start to have high reliabilities and therefore mostly positive expected utility for the society too and because it is less careful, it gets more contracts. The best negative fee policies remain close to the *Constant 0.0* policies, but the statistically significant difference between them remains with the *Very Large* and *Huge* seller populations (at the $p < 0.05$ level in *Very Large*, at the $p < 0.001$ level in *Huge* and at the $p < 0.0001$ level in *Large* seller population). In the *Reverse Exponential* distribution, most opponents will have very high reliabilities and therefore all the best policies

[10] Here and in other cases we used two-sided tests, because we might get differences in both directions.

[11] In a one-sided test, the difference is still significant.

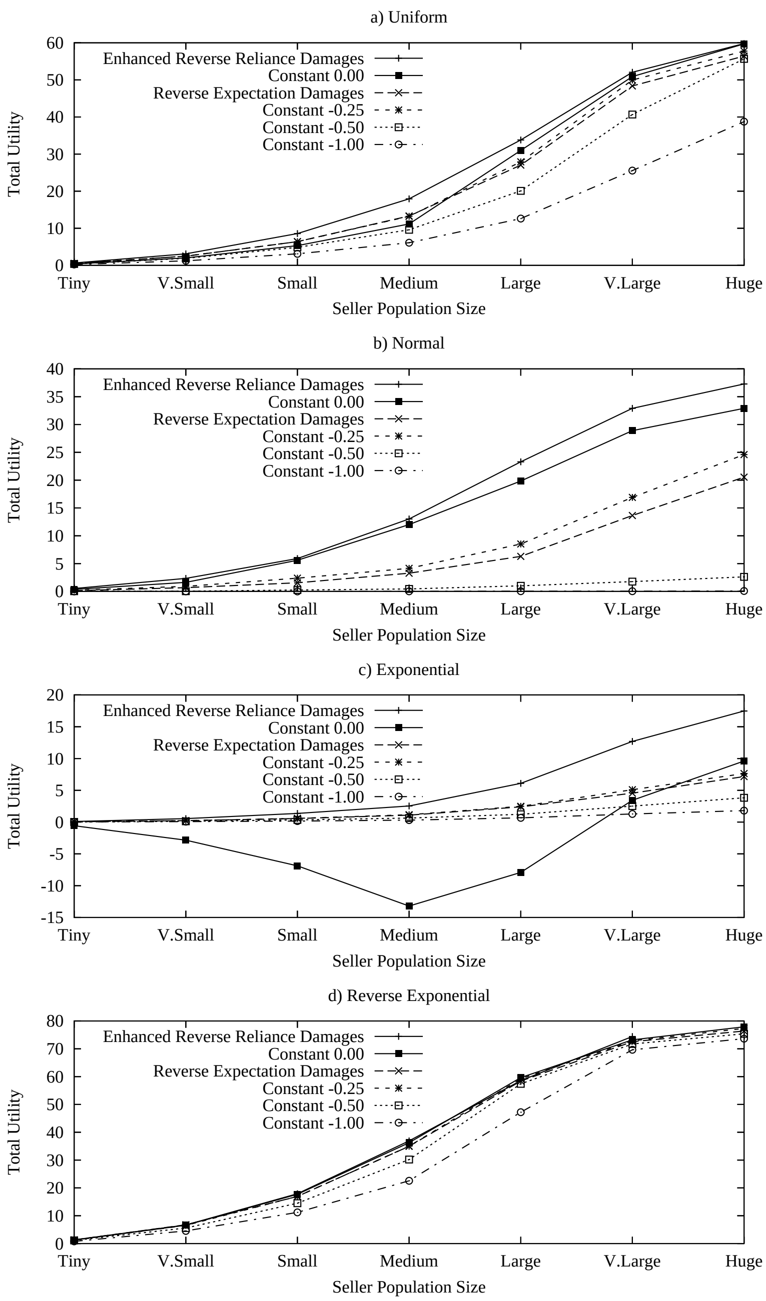

FIGURE 7.2: Policies with negative fees vs. Best policies (Hypothesis 15).

will perform quite well (no difference statistically). Only the policies with the largest negative fees are still too careful when choosing who to negotiate with and consequently they fare less well. Here, the *Constant 0.0* policy is not able to beat the best of the negative policies at all and it actually loses to it in *Small* and *Medium* seller populations ($p < 0.05$ and $p < 0.01$ respectively). This may be because even in the *Reverse Exponential* policy some providers with mediocre or even low reliability exist in a *Small* or *Medium* seller populations (not necessarily in the smaller ones) and when the policy is *Constant 0.0*, the agent will negotiate with these bad providers and that hurts the common good. In contrast, when the best policies with negative fees are used, such providers will often be ignored. When the *Constant 0.0* policy is able to choose more reliable providers from a bigger population, the difference between the policies vanishes.

7.3 Summary

In this chapter, we discussed selection decision and investigated how the decommitment policy affects the buyers' choice between different contract partners, who have a verying level of reliance as well as quality and how these choices affect the common good in a variety of settings. We investigated several near-optimal policies and compared their performance with the optimal and other policies (contribution **C4**).

In more detail, compensating for the buyer's losses to any extent in a case where the buyer chooses between heterogeneously reliable sellers, is problematic from the society's point of view. This is because any compensation in case of non-performance, makes that possibility less worrying and encourages the buyer to choose less reliable providers and, with fees higher than the expected utility, even targetting especially the unreliable sellers. The best decommitment policy from the selection decision's point of view would be to make the buyer pay the seller's costs in case of non-performance, even though it is the seller that is forced to decommit. However, this policy is not without its problems, and some of them are quite devastating. First of all, the both *Reverse Reliance Damages* policies assume that the buyer is the only one making the decisions and that the seller's reliability is fixed. Now, let us assume that we use the *Enhanced Reverse Reliance Damages* policy, have selected an opponent, negotiated with it and reached an agreement. Now the seller has to make a decision about whether or not to stay in the contract (performance decision). Its utilities are:

$$\begin{aligned} U_s(q \mid \text{performance}) = &\ \ p - C_s(q) \\ U_s(q \mid \text{decommitment}) = &\ \ -f - D(t)C_s(q) \end{aligned}$$

The fee covers the possible cost $D(t)C_s(q)$ and the 'seller's half' of the profit or $\frac{1}{2}(V_b(q) - C_s(q))$, so we get:

$$\begin{aligned} U_s(q \mid \text{performance}) &= \; p - C_s(q) \\ U_s(q \mid \text{decommitment}) &= \; \tfrac{1}{2}(V_b(q) - C_s(q)). \end{aligned}$$

This means that if the final price is lower than its average profit, the seller will actually prefer decommitment even without any adverse effects! A similar situation could also arise with all the *Constant* policies with negative fees.[12] And these contracts would still be beneficial for the society, so this is hardly the situation we were hoping for. In the experiments described in this section, the sellers were only allowed to decommit if they experienced an adverse effect (i.e. because of their reliability) and not for this reason. The performance of the *Enhanced Reverse Reliance Damages* policy would not be anywhere near as good as it is now if this restriction was not in place. However, here we were interested in the selection decision and the selection decision alone and in that setting, it is the optimal policy. The *Reverse Reliance Damages* or *Constant 0.0* policies would be better in this respect, because the former would give $U_s(q \mid \text{decommitment}) = 0$ and the latter $U_s(q \mid \text{decommitment}) = -D(t)C_s(q)$, both of which are less than $p - C_s(q) > 0$. Therefore these policies would never mean decommitment without some sort of adverse effect.

[12]When the fee in the *Constant* policy is very high, this is almost a certainty. The seller's profit can be at most 0.5 and that would require a price of 1.0, which is not very likely. Only a buyer with $q_b^{max} = 1$ would be willing to pay that and such buyers are not very probable. So, basically, if the fee is -0.50, -0.75 or -1.00, the seller would *always* choose to decommit. The only thing stopping it might be that decommitting always would eventually decrease the reliability rating and less buyers would choose to negotiate with it, but that might take a while to happen.

Part II

Concurrent Bilateral Negotiation Strategies

Chapter 8

An Adaptive Concurrent Bilateral Negotiation Model

After discussing the market setting and the common good in the previous part, we will now here (in part II) discuss a rather more complicated setting and concentrate on the buyer agent's strategies and welfare alone. This allows us to try more complicated strategies and tactics and see more clearly the effect they will have on the individual's performance. We will do so in a somewhat more complicated setting, namely concurrent bilateral negotiation, where one buyer agent can negotiate with many provider agents at the same time. Nevertheless, the results and ideas from the previous part are still very much in use here. In particular, we still have decommitments and decommitment policies and we see how the buyer agent should adapt to the possibility of changing circumstances. Here, the changing circumstances may force buyers and/or sellers to decommit and the buyer agent should take this into account when deciding what to do.

In more detail, concurrent bilateral negotiation is a more complicated form of interaction than the market setting we discussed in the first part. Thus, more factors need to be considered. Of the decisions discussed in the first part, we will concentrate here on the contract and selection decisions, deciding when and if to enter into a contract and who to negotiate with. We focus on these two because we want to have only a couple of decisions and given this, we want to discuss the decommitment policies' effect before and during the negotiation.[1] This is because after the contract has been formed the environment loses most of its significance[2]

[1]The other two decisions (performance and reliance) occur after the contract has been formed.

[2]Obviously the possibility of re-entry and the probability of a replacement contract do matter as we discussed in chapter 4.

and only the contract and its possible decommitment matter. Also these decisions have not been discussed together in the literature.

In addition, we also consider issues like choosing a good tactic in a single negotiation and coordinating negotiations over the same and even different services. To this end, this chapter provides an overview of our adaptive concurrent bilateral negotiation model and the remaining chapters of this part (chapters $9-11$) will discuss the details and our experimental results. The work in this part is related to the research contributions **C5**–**C9**. This chapter is concerned with the contribution **C5**, designing a model for concurrent bilateral nexgotiation, the rest of the contributions will be discussed in subsequent chapters.

We start the introduction of our model by explaining the environment the buyer agent needs to be able to work in (section 8.1). We will then discuss the model itself (section 8.2) and its implementation (section 8.3). Finally, we will summarise the main points (section 8.4).

8.1 The Marketplace

We start by discussing the environment or the marketplace. In most respects, the marketplace is just a variation of the basic market model of the previous part. However, we have also made some changes for the reasons we will discuss here and later.

So, we have a marketplace where buyers and sellers meet to negotiate on the price of services. Unlike in the market setting, however, we can also have multiple separate markets offering different services, although in most cases we only investigate one market at the time. We also only have one buyer (subscript b) and many sellers (subscript s) in each market,[3] although the buyer may negotiate with more than one seller concurrently in any or all of these markets. Since in this chapter we are developing the buyer agent for such a setting, we are interested mostly in the buyer's utility, U_b. Otherwise the basics are the same as in the market model (see chapter 3). As before, the time t is discrete and divided into turns. We assume that all participants expect the delivery of the service to occur at the same time $t_{delivery}$.

We will now explain how the marketplace works (section 8.1.1), then we discuss how the negotiations proceed and how the parties get the parameters for the

[3] And it is the same buyer in all the markets.

negotiations (section 8.1.2). Finally, we will discuss how we make the parties consider the issues associated with decommitment on the contracts that have been formed (section 8.1.3).

8.1.1 Matching and Entries

We still have the matching occuring every 100 turns and parties still negotiate for 100 turns on the price of the service. However, unlike in the market setting, we allow the buyer agent to choose the specific sellers it wants to negotiate with (opponent selection). In particular, we will experiment with several different types of opponent selection methods. The simplest one is, of course, the *random* selection, where the buyer chooses its negotiation partners among the sellers in the market at random. However, we will also discuss more sophisticated ways to make this choice.

Another significant change we have introduced is that we have diversified the negotiation tactics that the parties employ. In the market setting, we used only the simple *Exponential Time-Dependent* negotiation tactic. However, finding a good countertactic to that is relatively simple (especially if and when you have good information about the opponent's valuation and/or costs and hence, their reservation price or their deadline), because it would simply be a matter of waiting for the best offer. Thus, we will assume that the parties know each other's reservation prices, but not the deadlines. This assumption, although quite bold, is not unreasonable in many markets where the production technology and markets are well-known.[4] This means the sellers must be able to use other negotiation tactics as well (otherwise they would be left with almost no utility).

As discussed above, we want to be able to make interesting decisions already in a single negotiation, so the opponent will have to have different and more challenging negotiation tactics than was the case in Part I. However, on the other hand, we want the number of tactics to remain relatively small to keep the bilateral negotiation part manageable. To balance this tradeoff, recall that our main interest in this work is on the higher levels, so we endow the sellers with four different tactics:

[4] We could have introduced uncertainty about the opponent's reservation price but that is not likely to significantly change the situation and would be an additional complexity to our model. This means the performance we get from our model may be slightly too high for some settings, but such uncertainty would be easy to add to our model. We will discuss this aspect further in the future work (section 12.2.2).

- *Exponential Time-Dependent*: as in the market setting, the β parameter is still chosen at random,
- *Random* tactic: the offer is selected at random between the seller's reservation price and one (from a *Uniform* distribution). This makes it very hard to guess what a single offer might look like and it is unlikely that an actual reservation offer will be made. This means that simply waiting it out does not work.
- *Pure Behavioural*: the seller's offer mirrors the offer the buyer makes. If the buyer offers the seller's reservation price, the seller will offer the buyer's reservation price (if higher than its reservation price) and any improvement on that on the buyer's behalf will be met with equal improvement on the seller's behalf. This basically means that the best contract for the buyer is the midway point between the reservation offers leaving half of the utility to the seller.
- *Random Behavioural*: the seller's offer mirrors the offer the buyer makes, but not accurately. Instead a random offer is chosen between one that would represent double the concession the buyer has made (between the seller's and buyer's reservation prices) and one. This combines the two previous approaches and means that one needs to consider carefully what offer to make.

The three new tactics mean that waiting it out making insignificant offers is no longer a viable tactic. With the *Random* tactic, the seller's offers go up and down and the buyer needs to choose when to accept one of them. With the behavioural tactics (the last two), it is no longer viable to keep making offers that will never get accepted until the seller makes a good offer. Rather the buyer is forced to make offers that the seller might accept to lure it to make a good offer. Now, as it turns out, there is an optimal counter tactic for each of these but they are, of course, quite different from each other. We will discuss these tactics and their counter-tactics in detail later (see section 9.1.2.2).

In more detail, the tactic (θ_s) is selected for each individual seller at random, all four usually being equally likely (so the probability of each is 0.25). The buyer agent has an a priori (without entering into negotiation) guess what the chosen tactic of a given opponent is with probability χ, otherwise it will have no information. Such guesses could be based, for example, on the previous encounters. The buyer's guess is correct with probability γ and these probabilities are (accurately) known by the buyer. The guesses may be a result of observations about the seller's preferences in earlier negotiations between the parties or they can be derived from some other information. If it always has a guess and its guess is always right ($\chi = 1.0$ and $\gamma = 1.0$), it will always know the opponent's tactic accurately, but this will not always be the case. Given this, we will investigate a range of possible values for both parameters and their effect on the buyer's performance.

We assume that the sellers take a failed negotiation as a sign that the buyer will not want their service and because there are no other buyers in the market, they always exit the market. Therefore, the buyer will never be able to negotiate with the same provider more than once. Instead, the buyer chooses new opponents to negotiate with from the remaining providers. If the buyer gets a contract, it will exit the market and there will be no further negotiations. Otherwise, the matching process is repeated 9 more times. This means that we will have 10 matchings in total like we did in the market setting. Also the entries and exits occur in the seller population as they did in the market setting. The contracts are performed when the negotiations end, $t_{delivery} = 1000$.

In the beginning ($t = t_0 = 0$), there are n_0 sellers in the market. Over time, some sellers enter and some may exit. As in the market setting, the numbers of entries for the parties are independent variables and follow the standard Poisson distribution, with the parameter $\lambda(t) = i\frac{t_{lastEntry}-t}{t_{lastEntry}}$, where i is the basic entry intensity, $t_{lastEntry}$ the last turn that entries are possible and t is the current turn.[5] Here, we only use *Medium* population because it gives us a reasonable amount of sellers (around 250) and trying different population sizes was not considered a priority in this work (much of the results discussed would not be affected by population sizes) so, $n_0 = 50$ and $i = 0.4$.

8.1.2 The Negotiation Parameters

As in the market setting (section 3.2), we still assume that the provisioning of the service costs money for the buyer. Specifically, in order to provide the service at the delivery time, the provider s has to invest a cost c_s at time $t_{c,s}(< t_{delivery})$. Here we only use the simplest setting:

[5] As explained earlier, this formulation means that entries are more probable earlier in the experiment. This is realistic because the parties are more likely to find a contract if they enter early. This is especially true for the providers, because we assume that the provision of the service takes time, t_p, and they cannot wait until the last moment to find a consumer.

- *Any*: $Uniform(\max(0, t_{e,s}+1), 1000)$,

where $t_{e,s}$ is the time of entry for provider s. The time $t_{c,s}$ is selected independently for each provider using the same interval. Each provider has a quality q_s, which is selected at random from $Uniform(0, 1)$. The cost is a function of quality and time:

$$C_s(q_s, t) = \begin{cases} 0, & \text{if } t < t_{c,s}, \\ c_s = 0.5q_s, & \text{if } t \geq t_{c,s}. \end{cases}$$

These provider characteristics are mapped into typical bilateral negotiation parameters by setting the reservation price r_s equal to the provider's preparation cost c_s and the deadline to $t_{c,s}$. This means that the provider will never accept a price that is less than its costs and that if the provider does not have a contract when it should start preparing for service, it will exit the market. The provider's utility for a contract is: $U_s(p, c_s) = p - c_s$, where p is the contract price. The sellers do not consider decommitments when they make their decisions and they are, therefore, vulnerable to the buyer exploitation. By using simple sellers, we can see the effect of buyer strategies and tactics more clearly. In a more realistic setting, we might not be able to see the differences the buyer strategies make or at least they would be obscured by the sellers' reactions.[6]

The buyer does not have costs, but it has a deadline $t_{x,b} = t_{delivery} - 1 = 999$. The consumer's utility for the contract is $U_b(q, p) = V_b(q) - p$ and here we use a simplified version of the buyer's value function, $V_b(q) = q$. So, effectively we have removed the minimal acceptable quality which means that the buyer will be able to benefit from any provider's service. This change is purely technical and is done to remove an extra source of variation in our results.

8.1.3 Decommitment

In the market setting, we introduced an adverse impact on the contract utility that could affect one or both parties. When it occurs, it makes them want to consider decommitment. We have these impacts in here as well and, as before, we use a_b and a_s to denote the probability that the buyer or seller (respectively) will be affected.

For the provider, the decrease means that the cost of providing the service increases by amount L_s and this will decrease its utility by the same amount. He will then

[6]See our discussion and results on the contract decision in the market setting, chapter 6.

need to make a decision on whether or not to decommit from the contract in this new situation. The decision is influenced by the decommitment fee f_s. We assume that the provider s will decommit at turn t if and only if:

$$\begin{aligned} U_s(\text{contract}|L_s = l) &< U_s(t_{decommit} = t) \\ p - c_s - l &< -f_s - C_s(q_s, t). \end{aligned}$$

where $U_s(t_{decommit} = t)$ is the seller's utility, when he decommits at turn t and l is the amount the utility decreases. Here we use the following ten values $l \in \{0.1, 0.2, ..., 1.0\}$. So, the seller decommits if the decreased utility is lower than the cost it has already paid and the decommitment fee has to pay to get out of the contract. The seller learns of the loss at some point t_l (selected at random) between the time the contract was formed $t_{contract}$ and the time it was due to be performed ($t_{delivery}$) excluding both of the extremes. However, we assume that this loss itself is always avoidable, if the contract is abandoned before it is delivered. This means that the additional cost has to be paid just before the delivery. It is not possible that this additional cost is paid if there is no delivery.[7] There can only be one effect per party and the effect is always final. All possible moments for learning of the effect are equally likely, and unlike in the market setting, the effect can take place also before the entry, i.e. $t_l \sim Uniform(0, t_{delivery} - 1)$. A seller that is in the market when the effect occurs will immediately exit (withdrawing from a negotiation if necessary) and if the effect occurs before the entry, the buyer will not enter the market at all. This change was made to balance the exits between the buyer and the sellers. Otherwise there would have been more last minute exits for the sellers.[8]

The same applies to the buyer, except for two things. Unlike the seller agents, which can negotiate only with the buyer agent, the buyer agent can be negotiating with more than one opponent at the same time. This means that it may end up with more than one contract. Of course the buyer agent should be able to avoid accepting more than one contract itself, but if it makes offers to more than one opponent they may be accepted effectively at the same time. Getting into more contracts than required can be counter-productive because the buyer will then have to pay the decommitment fee for these extra contracts and these fees will usually decrease his utility.[9] On the other hand, negotiating with more than one

[7] However, it is possible that the seller has to pay the original cost even without a delivery as explained earlier. Only the extra cost is tied to the actual performance.

[8] This is because the buyer always enters at turn 0, but most sellers join later.

[9] We will of course also discuss the case where the fee is zero and therefore decommitting from extra contracts is free.

opponent at the same time will mean that the buyer agent is able to see a bigger part of the market and therefore, potentially, it will be able to find a better deal. This trade-off is of course one of the topics for discussion in this part and we will discuss various ways of making it (in section 10.1.2.2). Another difference is the costs. Like the market setting, we assume that the buyer has no costs associated with accepting and using the services (or more likely that these costs are built in its value function), so we always have $C_b = 0$ (for the seller $C_s \geq 0$).

In the market setting, the buyer's contract utility decrease had an upper limit, the adversely affected contract value could not be lower than -0.05 ($V_b(q) \geq -0.05$), because we assumed that the buyer can ignore the seller's performance at a low cost. However, this may not always be the case. The seller may, for example, need information or some other services from the consumer in order to produce its service and although this information or corresponding services may normally be very cheap (even 'free') to produce, an adverse impact may make this very expensive. For example, the buyer's Internet connection or hardware may experience a catastrophic failure and although it may be possible to get a new connection or hardware in time to provide the services the seller needs, this may be very expensive. In this case, the only way the buyer can avoid paying these unexpected extra costs is to decommit from the contract and that means paying the applicable decommitment fee. Here, we also assume that the adverse impact is always catastrophic or at least bad enough that the agents (both buyers and sellers) always decommit when it happens ($L_b = L_s = 5.0$). This assumption is made to keep the model more manageable and the results more clear.

Since, in this chapter, we are interested in the buyer utility. The buyer utility for the different possible outcomes are:

- *success*: $U_b(success) = V_b(q) - p = q - p - \Delta$
- *the buyer decommits*: $U_b(b\ decommits) = -f_b - \Delta$
- *the seller decommits*: $U_b(s\ decommits) = f_s - \Delta$

Δ denotes the decommitment fees that the buyer has had to pay to decommit from the extra contracts (those above 1). In theory, it is also possible that both parties decommit at the same time, but we did not consider that remote possibility in this work. One of the parties would always decommit first, even if both happened to decommit at the same turn.

8.2 Architecture for the Consumer Agent

In this section we discuss the architecture for our consumer agent. The agent's task is to find a service its master requires for a reasonable price. This basic architecture is based upon other models on concurrent bilateral negotiation, especially the model of Nguyen and Jennings (discussed in section 2.3.3.1). However, we extend and modify their model to also consider different negotiation tactics and interrelated negotiation groups. This makes the distribution of duties on different levels more explicit, clear and extendable. So what we have is a hierarchical model, where the different levels have a different view of the problem and the view broadens from the bottom up. The final decisions are taken at a high level so that a wider range of issues can be considered. However, the know-how and understanding of the lower levels are also utilised and the higher levels do not second-guess the recommendations from lower levels but adjust them only based on their wider view and in cooperation with the lower levels. Each level concentrates only on certain type of problems and assumes that other levels do their part. This makes the architecture easier to understand and implement and also allows more sophisticated solutions to single problems.

The consumer agent architecture consists of three layers or levels (see figure 8.1), which are:

- the `Negotiator`[10] level: Each `Negotiator` is engaged in a negotiation with a single provider agent over a single service. It only cares about that one negotiation and is unaware of anything else.
- the `Controller` level: A `Controller` manages a group of `Negotiator`s that negotiate about the provision of the same service. All `Negotiator`s under it report on their progress every turn. This gives this component an overall view of all negotiations on the same service, and allows it to manage the problems that concurrent bilateral negotiation can cause.
- the `Coordinator` level: The `Coordinator` manages the different `Controller`s and gets a report of their progress every turn. This gives it a good overall view of all the services and negotiations and enables it to manage dependencies between negotiations on different services.

[10] Throughout this thesis, we will use use the typeset text for the different levels of the model (for example the `Controller` (level)) and typeset text with surrounding inequalities (for example a `<Controller>` (component)) to refer to specific components. In addition, we will use the former to refer to the instances of a certain component (for example a `Controller` instance).

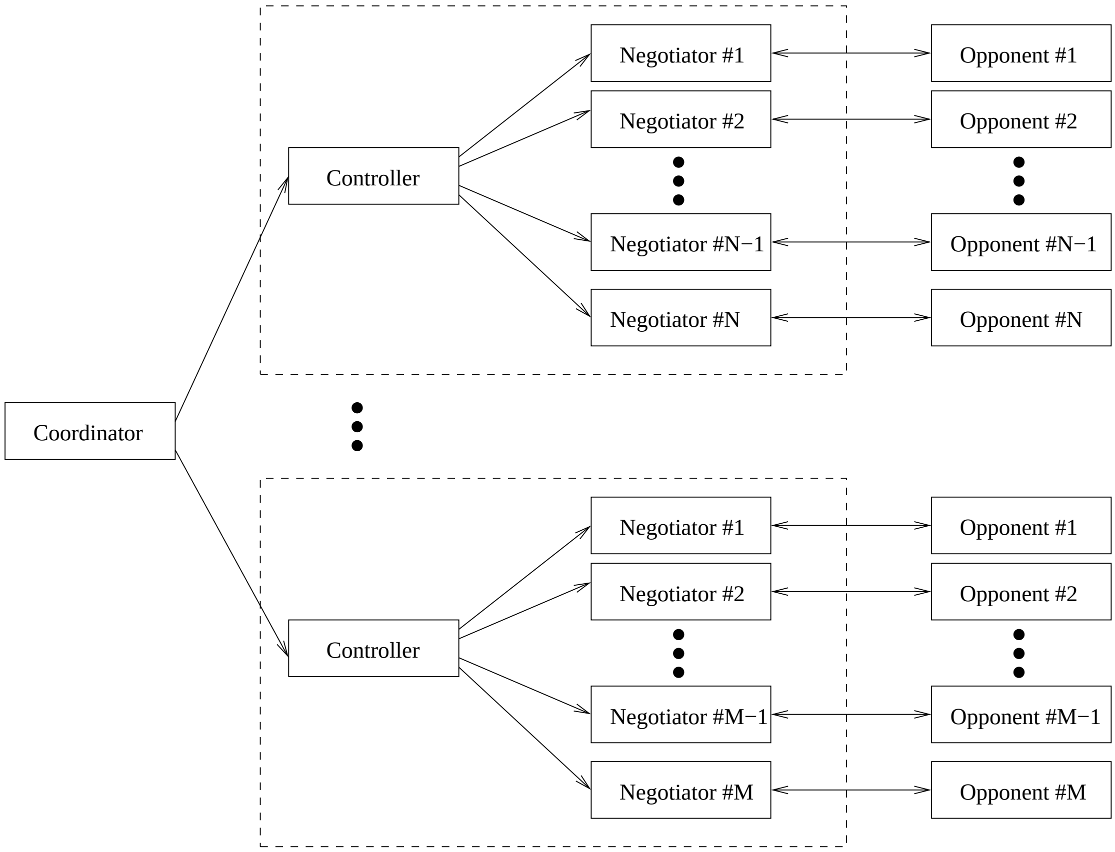

FIGURE 8.1: The Concurrent Negotiation Agent Architecture.

In more detail, the `Negotiator`s are trying to find the best offer to make in a negotiation with a given opponent. This is the `Negotiator`'s core competence, so the other levels will not change the offer to make without consulting the `Negotiator` first. The `Negotiator` will also suggest accepting an offer or withdrawing from its negotiation, but the final decisions are usually made at the upper levels.

Since we use concurrent bilateral negotiation as an interaction model, there is more than one `Negotiator` negotiating with the different providers on the same service in parallel. They have all the information that is relevant to this task and nothing more. They know all about possible opponent strategies and good counterstrategies, but know nothing of the other negotiations or the environment around them. Whatever information is considered necessary for the task at hand will be provided by the `Controller`. The `Controller` will also choose the opponent for the `Negotiator` (although it may consult the `Negotiator` during the selection process). The most important piece of information given to the `Negotiator` is the expected offers in future negotiations. This will allow the `Negotiator` to put the current negotiation in context and if a good enough result (better than something in the future) is not possible, it will know not to suggest an inferior deal.[11]

A group of `Negotiator`s is controlled by a `Controller`, which has two main tasks:

- opponent selection: choose the opponents to negotiate with and
- concurrency strategy: choose the number of negotiations to have at the same time.

These tasks are of course interrelated. The best number of negotiations may well depend on the opponents considered. These things may and often do depend on the negotiation strategies and available information about the opponents and that is where the `Negotiator` comes into play. The `Negotiator` will have the information about likely outcomes against different opponents and the best opponent selection and concurrency strategies use the information only the `Negotiator` can provide. However, the division of responsibilities is clear: the `Negotiator` can only advise on likely outcomes in a single negotiation, the `Controller`'s task is to consider the effect the multiple negotiations can have. That is, the risk of getting into more than one contract on one hand and finding the best providers on the other hand. Basically, a `Controller` makes all the decisions that affect more than one negotiation: it decides when it is time to quit a negotiation, when to accept a certain offer, when to start new negotiations and who to negotiate with. In the case where more than one of the agent's offers gets accepted by its opponents, a `Controller` decides which contract to take and which to decommit from.

A group of `Controller`s is, in turn, controlled by the `Coordinator`. The `Coordinator`'s task is to ensure that the interdependencies between the different services are satisfied. We will explore two types of interdependencies: substitutes and complements. *Substitutes* are services that although different essentially provide the same thing. Therefore, the buyer will only need one of the services. In case of substitutes, the `Coordinator`'s task is to manage the risk that the buyer gets more than one service and will need to decommit from the extra contracts by paying the decommitment fee. This is essentially very similar to the `Controller`'s problem and many of the solutions that work in one problem work in the other. *Complements*, on the other hand, are services that complement each other. Thus, having services A *and* B is better than having one or none of them. Here, we

[11] If getting that good a result is not possible in the current negotiation, the `Negotiator` will make sure that no contract is formed either by withdrawing or by making offers that it knows will not be successful.

investigate complements that have a positive value only if all complementary services are acquired and zero value if one or more of the services is missing. In such cases, the `Coordinator`'s task is to make sure that either it gets all the services or it gets none of them. This is because we again assume that the buyer will have to decommit from useless contracts and here, it means all contracts unless the `Coordinator` has a contract for all relevant services. This problem is slightly different from the others.

The interaction of these three components is described in figure 8.2. The basic idea is that the lower level makes a suggestion for an action (since it is an expert on its own level), but the decision is made at the higher level, which has a better overall view of the situation. The higher level usually follows the recommendation of the lower level, but it will sometimes adjust it (in cooperation with the lower level), if it knows of some conflict on its own level that the lower level was not aware of. So, when a `Negotiator` receives an offer from the opponent (step 1), it compiles a report of the negotiation so far and recommends an action to be taken to a `Controller` (step 2). A `Controller` waits until it has this information from all its `Negotiator`s and then devises an overall strategy taking into account the developments in all the negotiations. This recommended strategy with a status report is sent to the `Coordinator` (step 3). The `Coordinator` then creates an overall strategy over all negotiations taking into account the interdependencies between the different negotiation groups. This can mean starting new `Controller`s or ending existing ones. In addition, each `Controller` is sent its part of the strategy (step 4). The `Controller` then implements the strategy by starting new `Negotiator`s or sending existing ones their orders (step 5). Each `Negotiator` then follows its orders by taking the prescribed action (step 6).

Due to the layered structure, the `Coordinator` will not know how many `Negotiator`s there are in each market or what sort of offers they will be making. It does not need this information but trusts that the lower level to do their duties. It only needs to know the probability of success and expected outcome in each market and this information will be provided by its `Controller`s.

8.3 Implementation

The implementation was extended and modified from the marketplace model implementation we used in part I (section 3.4) using the same tools. The structure of implementation for the buyer agent follows the architecture described, each

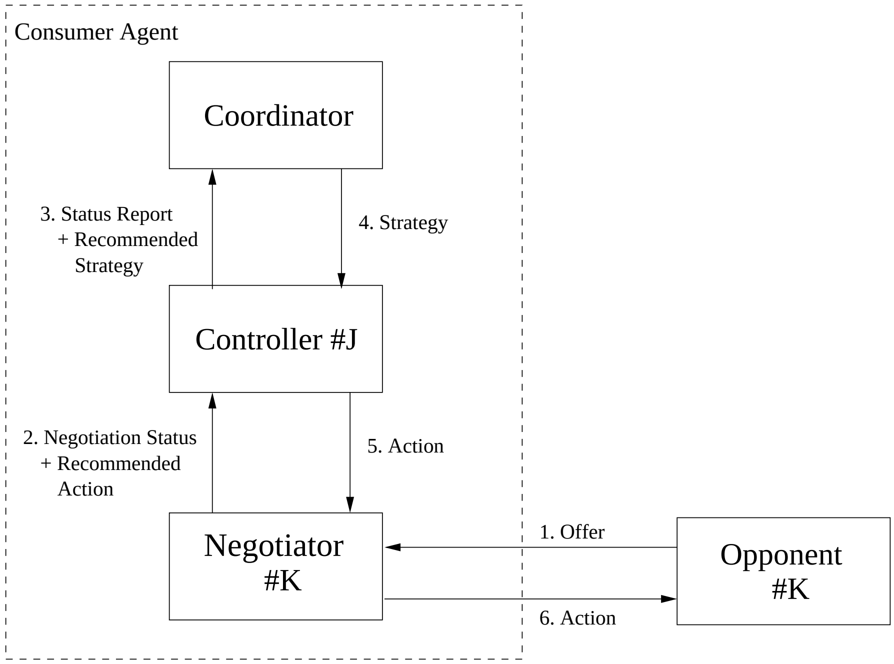

Figure 8.2: The interaction of the consumer agent components.

component is a class of its own and different strategies are implemented as classes with the same interface (that the other parts use to interact with them). Special care was taken to ensure modularity of the implementation, so that strategies and tactics can be changed at runtime. As in Part I, a configuration file contains the information about the strategies, tactics and other settings to be used and the system is configured according to the contents of the file at runtime. The file generated has 250-2500 values for each run depending on the settings used. The highest number of settings are in cases where the `Coordinator` and several markets are used, because events at each market are described in detail. The analysis is conducted in the same manner as in Part I.

8.4 Summary

In this chapter, we described the environment the buyer agent has to work in and we offered an overview of our model for the consumer agent engaged in concurrent bilateral negotiation. Specifically, we have a layered-model where each layer or level makes certain types of decisions in the concurrent negotiation. We have `Negotiator`s that handle single bilateral negotiations, we have `Controller`s that manage potentially many bilateral negotiations on the same service through managing a group of `Negotiators` and a `Coordinator` that manages negotiations on different services and the interconnections between them. Each different question is isolated to a separate and interchangeable module and we defined the interaction between these modules in detail (contribution **C5**). The next step is to discuss these levels one by one in more detail and test our ideas empirically. This will be done in the subsequent chapters $9 - 11$.

Chapter 9

Negotiation Tactics: The `Negotiator` Level

As explained in section 8.2, the tasks of a `Negotiator` include:

- *Managing one bilateral negotiation*: A `Negotiator` is responsible for exchanging offers with its opponent and reporting results of these exchanges to the `Controller` every time something happens.
- *Optimising negotiation offers*: A `Negotiator` is responsible for deciding what offers to make in a negotiation. In order to do this, a `Negotiator` will have to consider the instructions from the `Controller`, the opponent's characteristics (especially its quality and the estimate of the negotiation tactic the opponent will use) and the possibility (if any) of one or both parties needing to decommit.
- *Providing estimates on the negotiation outcomes and success probabilities*: To function efficiently, the higher level may need an outcome estimate of any on-going or considered negotiation. This estimate requires both the expected utility of a successful outcome and the probability of achieving that result. Providing this information for any opponent at any time (before or during the negotiation) is the `Negotiator`'s task.

Starting the negotiations, selecting opponents, considering other negotiations, setting targets and everything else is done by the the higher levels, usually the `Controller`. Also the actual negotiation tactic (or the method for finding the best offer to make) is given by the user and the `Negotiator`'s optimisation task is to use the given tactic as effectively as possible in the circumstances. Our aim in this chapter is to discuss our contributions **C6** (bilateral negotiations and the Negotiator) and, in part, **C9** (future offers).

Given this background, we will first explain the architecture of the Negotiator level and its interaction with other parts in more detail (section 9.1). We will then employ the Negotiator in the series of bilateral negotiations and discuss these experiments and their results 9.2. We will conclude this chapter with a summary of our findings (section 9.3).

9.1 Architecture of the Negotiator

We discuss the architecture in three parts. First, we give an overview of the architecture and how the Negotiator interacts with its environment (section 9.1.1). Second, we will discuss specific negotiation tactics that the Negotiator (and its opponents) use in detail (section 9.1.2). Finally, we will discuss how the Negotiator can, with the help of a Controller, take into account the offers that the agent might receive in the later negotiations (section 9.1.3).

9.1.1 Overview

The <Negotiator> component consists of two major parts:

- <NegotiatorBase> module deals with communication with both the <Controller> and the opponent with respect to all associated synchronisation issues. This functionality is needed no matter what tactic a Negotiator uses.
- <Tactic> module takes all tactic-specific decisions. Each tactic or family of tactics will have its own <Tactic> module, but they all implement the same simple interface for the <NegotiatorBase> to use. This makes different <Tactic> modules interchangable and allows the selection and configuration of the appropriate tactic at run-time. In addition, this structure allows a complete change of tactics, even during an on-going negotiation (although we do not do that here) simply by changing the Tactic instance to another.[1]

[1] Since each such module is also configurable, smaller changes can be achieved by changing the configuration of the module. All tactics have three parameters in common: a deadline, a reservation price and an optimal price. Some tactics may also have additional parameters.

The basic operation of a `Negotiator` is described in figure 9.1. In particular, a `NegotiatorBase` module will get an offer from its opponent (step 1). It will then ask its `Tactic` module to provide an analysis of the negotiation status and recommend a course of action (steps 2 and 3). This information is given to the `Controller` (step 4). The `Controller` makes adjustments that are necessary due to the overall situation (for example other negotiations) that the `Negotiator` is unaware of and returns the action to be taken and possibly some changes to negotiation parameters (step 5). The `NegotiatorBase` will then inform its `Tactic` of these changes and take the prescribed action (step 6). In the experiments we discuss in this chapter, the `Controller` will simply approve whatever decision the `Negotiator` recommends.

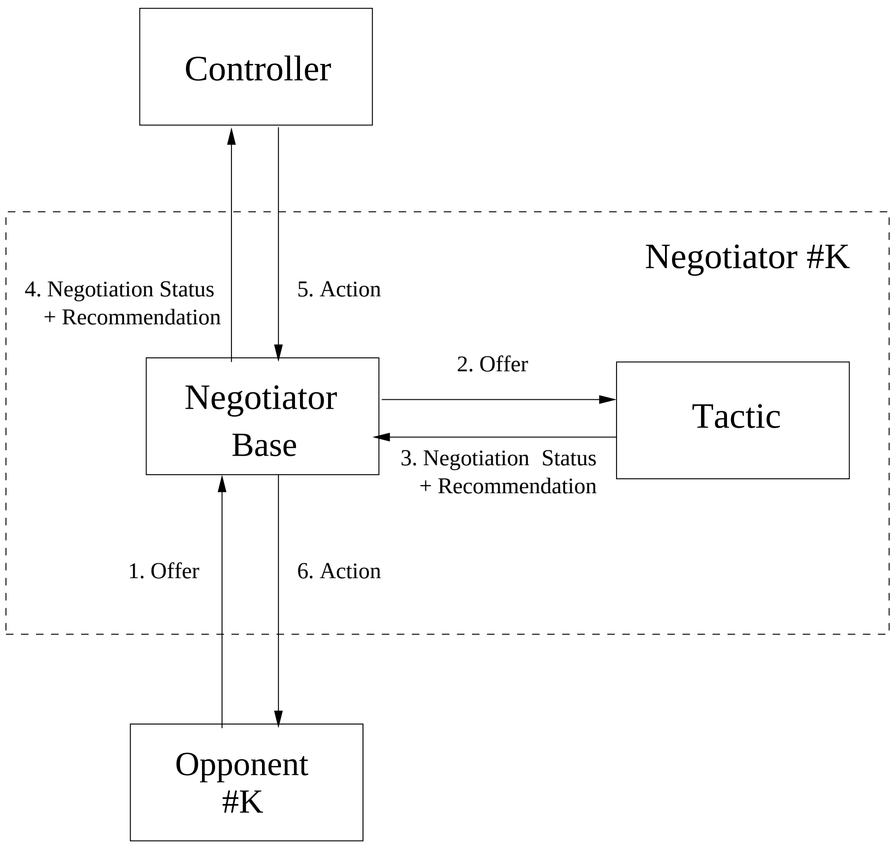

FIGURE 9.1: The basic operation of a `Negotiator`.

9.1.2 Negotiation Tactics

We will then take a closer look on the `<Tactic>` module, since it is clearly the most interesting part of the `<Negotiator>`. The tactics are discussed in three parts. First, we discuss the seller tactics (section 9.1.2.1) and then the tactics the buyer agent will use against them (section 9.1.2.2). We conclude by discussing the most advanced tactic we employ, the *Adaptive Counter* tactic (section 9.1.2.3), which is a combination of all the countertactics.

9.1.2.1 The Seller Tactics

As explained in section 8.1.2, the sellers have four tactics at their disposal:

- *Exponential Time-Dependent*,
- *Random*,
- *Pure Behavioural*, and
- *Random Behavioural*.

We will now explain the three new tactics in more detail. The *Random* tactic does not really need that much more explanation. The seller using the *Random* tactic simply makes offers that are selected at random from *Uniform*($r_s, 1$), so a uniform distribution between the seller's reservation price and the highest allowed price (one). Every offer is drawn from the same distribution and the offers are independent of each other.

The behavioural tactics need more explanation. They use a concept of *concession* to determine their offer. Here, the concession refers to the amount of utility the buyer is willing to give to the seller with its offer. This is calculated as a proportion of total utility (value−cost), so the buyer's concession (α_b) is equal to:

$$\alpha_b = \begin{cases} -\frac{c_s - o_b}{c_s}, & \text{if } o_b < c_s, \\ \frac{o_b - c_s}{V_b(q) - c_s}, & \text{if } c_s \leq o_b \leq V_b(q), \\ 1, & \text{if } o_b > V_b(q), \end{cases}$$

where o_b is the offer made by the buyer. In the *Pure Behavioural* tactic, the seller will repeat the concession in its own offer o_s, so it makes an offer of:

$$o_s = \begin{cases} V_b(q) - \alpha_b(V_b(q) - c_s), & \text{if } \alpha_b \geq 0, \\ V_b(q) - \alpha_b(1 - V_b(q)), & \text{otherwise,} \end{cases}$$

So, if the buyer makes an offer that fails to exceed the seller's reservation price, the seller responds by making an offer that is sure to exceed the buyer's and the buyer's offer of zero (the minimum offer) will be met with an offer of one (the maximum offer). As with all other tactics, if the offer the seller would make is lower or equal to the offer the buyer actually made, it will accept the buyer's offer.

The *Random Behavioural* tactic is a combination of the two other new tactics. It uses the concession as per the *Pure Behavioural* but selects the offer at random from *Uniform(*$V_b(q) - 2\alpha_b(V_b(q) - c_s), 1$*)*, where concession α_b is capped at 0.5, so any value over 0.5 is considered to be 0.5. And of course if the concession is negative, a random offer from *Uniform(*$V_b(q), 1$*)* is made.

If and when the buyer uses these tactics, they of course work as mirror images of what was just described. This means that, for example, the *Exponential Time-Dependent Tactic* makes increasing offers instead of decreasing ones but the β parameter is selected at random as it is with the sellers. In the *Random Tactic*, the buyer version chooses offers between $[0.0, r_b]$ instead of $[r_s, 1.0]$ and so on.

The problem with these tactics is that it is very difficult to estimate the outcome with some of them, when they are used against each other. For example, the outcome of two opponents using the *Exponential Time-Dependent Tactic* depends on the β parameters and deadlines and it can be almost anything between c_s and $V_b(q)$. This means that it would be impractical to make the `Negotiator` use any of these tactics and still expect it to be able to estimate the outcome and success probabilities with any reasonable accuracy. So, although the buyer agent can and in some experiments we will discuss later does use these tactics, something else is needed for the concurrent bilateral negotiation and that is exactly what we discuss next.

9.1.2.2 The Simple Counter Tactics

The buyer agent, in most cases, uses a novel counter tactic to each of the four possible seller tactics. Specifically, all countertactics calculate an optimal offer to make and repeat that offer until it is accepted or the negotiation ends. Therefore the outcome of a successful negotiation is always known in advance and the problem is only to assess the probability of success.

Although there are only four different seller tactics and these counter tactics are specifically aimed against one of them at the time, they are more general than they may seem at first glance. They do not care about single offers but only about the distribution of the best (lowest) offer and that is calculated using information about the buyer's value and seller's costs and assumed negotiation tactic.[2] Although these counter tactics assume knowledge of the buyer's value and seller's costs, they would probably work reasonably well even if these parameters were estimated instead of known, although we will not address that possibility in this work.

So, each seller tactic has its own counter tactic. The *Exponential Time-Dependent Counter* tactic is the simplest of the counter tactics. It basically makes an offer that is equal to the seller's expected reservation price and keeps making that offer throughout the negotiation. If (and when in our experiments) the estimate is correct, an opponent using the *Exponential Time-Dependent* tactic (or any tactic that will at some point offer the reservation price) will eventually make that offer. In our setting, this optimal offer is basically $C_s(q) + 0.00001$, because the seller will want a positive utility and this offer will give it one (marginally).[3]

The second easiest counter tactic is the *Pure Behavioural Counter* tactic. Because we know that any concession we make over the seller's costs will be met with a similar concession under our value, we can easily deduce that the optimal offer to make is the half-way point. Any lower offer will not be accepted and any higher offer will give the opponent extra utility (that is away from us[4]). We make this offer at every possibility. If the opponent is using a *Pure Behavioural* tactic and we have estimated its costs (and it has estimated our value) correctly, we will have a contract in the first round of the negotiation: We make an offer that they will accept. The useful property of this tactic is that it works against pretty much any conceivable tactic that does try to get to a reasonable outcome. If the opponent at any point makes an offer that is lower or equal to the half-way point between value and cost, we will get a contract. Against some tactics, such as a very bullish (boulware) *Exponential Time-Dependent Tactic*, it may take a while (until the deadline in fact), but even the *Random* tactic is very likely to make that sort of offer at some point. The downside of this tactic is naturally that most tactics would make lower offers at some point and this tactic misses the opportunity to take such offers.

[2] *When* that minimum offer is actually made, is not relevant to these tactics as such, although `Negotiator`s are sometimes called to estimate the contract time as well and that requires some idea when that is going to be. We will discuss this later in chapter 10.1.2.2.

[3] We use offers that are rounded to the five decimal points so this is the smallest possible positive utility.

[4] A price negotiation is a zero-sum game. That is, one party's gain is always equal to the other party's loss.

The two remaining counter tactics are more complex. The problem is that in theory these tactics, especially the *Random Tactic*, can make any sort of offer and the trick is to set the offer so that the product of the probability of success (offer accepted) and the expected utility in case of success are maximised. The lower offers usually mean a higher utility but a lower probability of success and high offers mean low utility and a high probability of success. The tradeoff between the two is central. So basically we want to make an offer that yields us the best expected utility:

$$EU = P(\text{contract})EU(\text{contract})$$

where $P(\text{contract})$ stands for the probability of a contract and $EU(\text{contract})$ for the expected utility of that contract.

Now, if we only had one offer to make, the situation would be relatively straight-forward. For example, against a *Random* Tactic an offer of $\frac{1}{2}(1.0 + C_s(q))$ would have a 50% chance of getting accepted depending on whether or not the seller's random counteroffer would be higher or lower than that and since all offers are equally likely we get the mentioned probability. However, if we have a hundred turns of offers and counteroffers to consider, we can do better. The probability that all hundred offers would be greater than the said offer, is $0.5^{100} = 7.9 \cdot 10^{-31}$ so very small indeed. More generally, the probability of success against a *Random* tactic (that makes offers between $[c, 1]$) is:

$$P(\text{contract}|\text{random tactic}) = 1 - (1 - \frac{o - c}{1 - c})^t,$$

where o is the offer made, c the cost and t the number of negotiation turns. As long as the offer is higher than c and lower than the maximum offer 1, the term $1 - \frac{o-c}{1-c}$ will also be in the interval $(0, 1)$[5] which means that it is decreasing in t and therefore the probability of success increases when the number of negotiation turns increases. This means that the optimal offer is usually lower, sometimes even much lower, than what it would be if only one offer was made.

In the other random tactic, the *Random Behavioural* tactic, things are also more complicated. First of all, the seller's offer range depends on the buyer's offer and is not very straight-forward to begin with. We can ignore the cases where our offer is lower than the seller's costs, because that will never get accepted, and also the cases where we make a greater offer than $\frac{1}{2}(V_b(q) - C_s(q)) + C_s(q)$, because this

[5] If the offer is equal to c or 1, the t will have no effect on the probability, but it will be always 0 or 1, respectively. We do not consider any offers that are higher than 1, because such offers are not allowed.

offer will give us the maximum concession anyway. We get:

$$P(\text{contract}|\text{random behavioural tactic}) = 1 - (1 - \frac{o - v + 2(o - c)}{1 - v + 2(o - c)})^t$$

We could use these probability functions and formulae to calculate the contract's expected utility to determine the optimal offer analytically but this could be quite complicated, especially if the decommitment fees are compensatory and depend on the contract price too. Also it would not generalise to cases where the exact functions are not known but we have enough data to have empirical estimates for these probabilities. We have therefore decided to use another approach, namely going through all the relevant options. It basically involves going through possible offers we could make and calculating the expected utility for each offer and then choosing the offer that yields the best expected utility. Although this may seem like a lot of work, there are some very useful ways to optimise this process and make it quite fast. In addition, this approach has an advantage that it can be generalised to any situation where the success probabilities for different offers are known and this also includes cases where the probabilities are known empirically and not analytically[6] and when they fail to follow any simple formulae. Also, as we will shortly see, this approach is readily extendable to cases where there are more than one possible tactic or even more than one negotiation.

The optimisation is really very simple. Since the probability of success is increasing and the buyer's utility is decreasing over offer o, it seems reasonable to assume that the product of these two (both in the interval $[0, 1]$) reaches a maximum at some point between c_s and V_b. This is perfectly consistent with empirical results. When the offer is c_s, the probability of success is zero but the utility is very high. On the other end of the spectrum, if the offer is V_s, the chances of success are high but the utility is zero. When the offer moves from these end points towards the other end point, the expected utility increases and it peaks somewhere between the two. Many nearby values may be very close to each other and indeed even produce the same expected utility but there never is more than one peak. So the point of the optimisation is to find the peak. This is relatively straight-forward. We start from the lowest possible offer and calculate the expected utility and increase the offer by 0.0001, calculate the expected utility again and so on until the expected utility is lower than it was before. That means that the peak is somewhere between $(x - 0.0002, x)$, where x is the last offer we tried. We then go

[6] This means that we have a good idea what sort of offers the opponent makes or accepts but we do not know the reasoning (negotiation tactic) behind these actions.

through these options, starting with $x - 0.0001$ and moving upwards by 0.00001 increments. If the value at $x - 0.00009$ is higher than at $x - 0.00001$, we know that the peak is between $[x - 0.00009, x)$ and proceed until the expected utility decreases. If the value at $x - 0.00009$ is smaller than at $x - 0.0001$ we need to go in the other direction and proceed until the utility starts to decrease. We choose the starting points carefully.

In the *Random Counter* tactic, we start from the seller's reservation price, because the optimal offer is usually quite close to that. In the *Random Behavioural Counter* tactic, we use a different starting value, equal to $\frac{1}{3}(V_b(q) + 2 \cdot C_s(q))$ because any concession we make is met by a random concession that has a minimum at double the concession we made. The start value means that the buyer makes a 1/3 concession and because the minimum of the seller's counter offer will be double that 2/3 concession, they actually refer to the same point. Any lower offer would mean that there was no overlap and therefore no chance of success. These start points mean that finding the peak is usually very quick and takes only a handful of iterations.

Another point worth making here is that although in all counter tactics, we calculate the optimal offer and keep making that same offer throughout the negotiation, this may not be an optimal approach in all cases. With non-behavioural tactics (such as the *Exponential Time-Dependent* or *Random* tactics), we could simply keep making an offer of zero (that the seller will never accept) and use the optimal offer calculated only as a threshold to tell us when we should accept an offer. This would give the seller two advantages over making reasonable offers. First, it would always be the seller who decides when a contract is formed. This would be very useful when we consider having several negotiations concurrently and want to avoid extra contracts. Second, in the case of the *Random* tactic we could sometimes get slightly better deals. If the seller is using the *Random* tactic, it will accept the buyer's offer if his next offer would be equal or less than the buyer's. If the buyer does not make non-zero offers, the seller would make that offer instead and sometimes that offer would be lower than the buyer's threshold offer.

However, this does sound more than slightly one-sided and it would also make concurrent negotiation too easy for our agent, so we assume that if the buyer does not ever make offers that are equal to these optimal offers (or higher), the seller will not make any reasonable offers either. So, even if the tactic in general is non-behavioural, we assume that there is this small behavioural aspect to it and, therefore, the buyers are forced to make these optimal offers and our buyer agents always do. Of course with the real behavioural tactics (*Random Behavioural* and *Pure Behavioural* tactics) this is not an issue because these tactics do force the buyer to make reasonable offers if it wants to have any chance of success.

The problem with these tactics is that they work very well against their target tactic, but may be less efficient against other tactics. This inefficiency comes in two forms:

- *Failure to succeed*: The negotiation will usually be unsuccessful. The optimal offer made by the counter tactic is too low to have a reasonable (any) chance of being accepted.
- *Failure to exploit*: The negotiation will succeed but the seller is given more utility than would have been strictly necessary had the right counter tactic been employed.

The *Exponential Time-Dependent Counter* tactic does very badly against any tactic except *Exponential Time-Dependent*, because its offer will not be acceptable to a seller using any of the other tactics.[7] It will therefore usually fail to succeed. On the other hand, the buyer using the *Pure Behavioural Counter* tactic can expect to get into a contract with almost all opponents, because its offer is acceptable to opponents using other tactics (in the case of *Random* and *Random Behavioural* tactics, this of course means just a very high (but still < 1) success probability) and it can take a while against an *Exponential Time-Dependent* tactic, for example.

These considerations are very important in a one-shot negotiation. However, if the same counter tactic is used against more than one opponent, the situation changes. For example, if there are 10 matchings, the `Negotiator` will get 10 chances to find an opponent that uses the tactic that is more suitable. Here, 9 failures may not mean anything, if you get that one success. Even the *Exponential Time-Dependent Counter* tactic may therefore work well if it has a reasonable chance of finding an opponent that actually uses the *Exponential Time-Dependent* tactic and the more there are opponents using that tactic in the population and the more opponents it tries to negotiate with, the better the chances. For example, in a population where 10% of the sellers use it, the buyer using this counter tactic will fail 90% of the time in a single negotiation, but if it can try 10 different providers, the chance of failure

[7] It might get lucky with the *Random* tactic but that is very unlikely.

decreases to under 35%.[8] However, with 25% using this tactic, the probability of failure in 10 attempts drops to 5.6%. And since the *Exponential Time-Dependent Counter* tactic does well on the exploitation part when successful, it may work very well in such a setting.

Another good thing about these counter tactics is that they can work blind. They do not have to know what tactic the opponent is using. They can just assume that the opponent is using the target tactic and calculate the optimal offer accordingly. If the opponent uses some other tactic, either the negotiation fails or you leave some utility to the opponent that could have been taken with a more appropriate tactic. Here, knowing the opponent's tactic only tells us whether or not there is any point in negotiating but it does not affect the outcome. So if the buyer agent doesn't know what tactic the opponent is using, it can always just try to negotiate and see if it works out or not. If it is successful, it knows what the expected utility is going to be because we keep making the same offer over and over again. If the probability distribution over different tactics (or probability that the opponent is using a tactic that will accept the offer the buyer is making) is known, the `Negotiator` can even give an estimate for the probability of success, if the higher levels need one. If there is a good chance of succeeding in the end, knowing what tactic a given opponent is using is not overly important. Of course in some settings, you get a limited number of chances to enter a contract and using a tactic that works against only some opponents may not be a great idea. We, therefore, introduce another tactic that can work against any opponent tactic.

9.1.2.3 Adaptive Counter Tactic

The *Adaptive Counter* tactic is a combination of all four counter tactics discussed in the previous subsection. There is no reason to limit oneself to these four, but the *Adaptive Counter* tactic could, in theory, handle any number of negotiation tactics. What it does need is the counter tactics, as per those described above, that can provide the probability of success for any offer made and a probability distribution over the opponent tactics.

This probability distribution is determined by the `Negotiator` itself, although the `Controller` provides it with two essential pieces of information:

- *its best guess on the tactic (and the reliability of this guess),*

[8]This is because $(1-0.1)^{10} = 0.3487$. We assume that the population is big enough so that negotiating with a few opponents in it does not alter the probabilities in a significant way.

- the *tactic distribution*: the probabilities for different tactics in the opponent population

In our experiments, the `Controller`'s best guess is either one of the four tactics or `unknown`, if no guess can be provided. The buyer will be provided with the reliability estimate γ, which gives the probability that the given guess is correct. This information can be used to generate the probability distributions over tactics. For example, let us assume that all tactics are equally probable in the seller population and that the `Controller` says that a given opponent is using *Exponential Time-Dependent* tactic and that we know any guess would be correct 40% of the time. The probability distribution is as follows:

$$
\begin{aligned}
&P(\text{Exponential Time-Dependent Tactic}) = \gamma = 0.4\\
&P(\text{Random Tactic}) = P(\text{Random Behavioural Tactic})\\
&= P(\text{Pure Behavioural Tactic}) = (1-\gamma)0.25/0.75 = 0.2
\end{aligned}
$$

All the other cases work in the similar way. The *Adaptive Counter* tactic does not try to recognise the opponent tactics while the negotiation is on-going but relies on these a priori probabilities instead.[9] Using these probabilities, the *Adaptive Counter* tactic calculates the optimal offer to make in a negotiation. If the opponent tactic is known (with probability 1), it will behave exactly like the appropriate counter tactic. However, it will be able to adapt also to incomplete information and probability distributions where more than 1 or even all four tactics have positive probabilities. It will use the same approach as the *Random Counter* and *Random Behavioural Counter* tactics, namely it tries all the relevant offers and chooses the one that gets it the highest expected utility. The price in a contract needs to be set so that the following evaluation is maximised:

$$EU = \sum_{t=t_1}^{t_4} P(\text{tactic } t)P(\text{contract}|\text{tactic } t)EU(\text{contract}),$$

where t_1 = 'Exponential Time-Dependent tactic', t_2 = 'Random tactic', t_3 ='Random Behavioural tactic' and t_4 ='Pure Behavioural tactic' and $P(\text{tactic } t)$ the probability that the opponent is using tactic t. Here the optimisation is slightly more challenging than in the single counter tactic case, but because the probabilities of success still increase in one direction and the utility of the contract decreases in

[9] This is because the four tactics we have would be relatively simple to distinguish from each other but this might not always be the case.

the same direction, it can be managed. We basically do the optimisation for each of the four tactics and choose the price that gives the highest expected utility.

The *Adaptive Counter* tactic tries very hard to adapt to the opponent's tactic, qualities and costs. This means that, with full information, it typically gets very high success probabilities against any opponent, unless there is a significant risk of requirement change and a high decommitment fee, in which case the *Adaptive Counter* tactic waits until an opponent can provide a positive expected utility and that may make it impossible to succeed in the early negotiations (because expected utility is typically negative). On the other hand, this also means that it does not usually wait for a more suitable opponent. Rather it tries its best to get a contract with the current opponent, no matter what its charateristics are. The *Adaptive Counter* tactic is as happy to reach a contract with a mediocre quality provider using a *Pure Behavioural* tactic than it is to enter into a contract with an excellent quality provider using *Exponential Time-Dependent* tactic. This is true even though the expected utility for the latter is much higher than for the former. This can be very useful if the buyer has a limited number of chances to find a contract, but in some other settings, it might be a problem. We will now discuss one, quite effective, method for curtailing this behaviour in settings where it might be harmful.

9.1.3 Considering Offers in the Future Negotiations

So far, the tactics have considered any contract that produces positive expected utility acceptable. However, this is often not the optimal approach, because we might later be able to find a better deal than what is currently on the table. This is because the probability of adverse effects decreases over time and because more options usually lead to better outcomes. We discuss these two reasons now in turn.

First, we consider the fact that *the probability of adverse effects decreases over time*. Since we assume that the adverse effects can take place at any time with the same probability, we can update the probability that we or our opponent will be affected, since we know that if we and they are still in the market, neither of us has been affected so far. This means that the expected utility of the same contract can either increase, decrease or stay the same over time. If an adverse impact is bad for the buyer (he is the one affected or the seller's decommitment policy will not cover the lost profit), the decreasing chance of adverse impact is good for the buyer and the buyer's expected utility increases over time. This effect is the stronger the higher the probability of adverse effect and the decommitment fee are. However, the adverse impact can also be beneficial for the buyer. This occurs when only the seller is affected and the seller's decommitment policy is over-compensatory (for example *Constant 1.0*). Here, the lower chance of adverse impact (and seller's decommitment) actually decreases the buyer's expected utility so the buyer's expected utility decreases over time. If only the seller can be affected and the decommitment policy is *Expectation Damages* (fully compensatory), the expected utility stays unchanged over time.

Second, *more options often mean better outcomes.* In any given negotiation, the buyer has a certain chance of getting a contract and it can achieve certain expected utility if it does. However, the success probability and the expected utility depend on the opponent's characteristics. If the buyer negotiates once with one opponent, all it can do is maximise its expected utility in that one negotiation no matter how good or bad the opponent is. Any contract giving a positive expected utility is better than no contract. However, if the buyer agent knows it is going to face, say, 9 more opponents after the first one, it can afford to be a lot pickier. So if the first opponent is abysmal, it can be ignored and maybe the next one will be better. There is of course a trade off here. A better opponent may not arrive after all or it might exit before an agreement can be reached and the current opponent may be the best one the buyer is going to meet this time in the market. However, this works also (especially) if there is no adverse impacts.

So the buyer agent has to know when to negotiate seriously and when to wait for a better opponent or improving circumstances. We will here describe one method for doing just that. The buyer's problem is described in figure 9.2. Basically at each matching (denoted by numbered boxes $1-10$), the buyer can either negotiate with a given opponent successfully (end up in a contract) or fail to do so (for whatever reason). In the former case, the buyer exits the market and in the latter, he will move on to the next matching hoping for better luck there. This is shown as arrows from each box. Now, if we assume that we know the quality distribution and the negotiation tactic distribution of the opponents, we can formulate a distribution for the future opponents and expected utilities they might provide.[10] However, this could become very complicated and would require quite detailed information that may not be available in a realistic setting. Moreover, our opponents vary in negotiation tactics, as well as in quality (two dimensional variation among the

[10]This is a type of *optimal stopping problem* (Ferguson 2008) where a certain action (negotiating with different opponents) is repeated and the problem is to decide whether the current choice is good enough to stop (get a contract) or if one should continue (go to the next opponent).

sellers). Moreover, we will later have many concurrent negotiations (even varying numbers), we can use different opponent selection methods, there are adverse impacts, and so on. This would mean that the analytic approach would sooner or later become very complicated and very difficult. We therefore use an alternative that uses empirical data gained from a given setting and then uses analytical tools to calculate the threshold utilities.

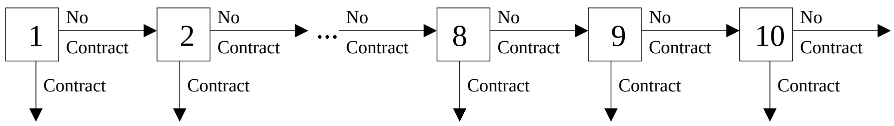

FIGURE 9.2: Structure of the Problem.

We basically run the market several (usually 5000) times and calculate the expected utilities in each matching as we see them. So if we have only one negotiation and we know we are going to negotiate with the first opponent on the list no matter what, we can simulate that by calculating the expected utility of the first opponent (taking note of its id) and then move on to the next matching, again taking the first opponent (but if we have already 'negotiated' with it, we will take the next one that we have not negotiated with) and so on. The best thing about this data gathering phase is that we do not have to actually start any negotiations. We can negotiate, of course, but since the counter tactics described above can provide us with an estimate of the outcome without negotiation, we do not have to negotiate to get this information. This can be useful, if we can use some time to observe the market before actually starting to negotiate in it.[11] In the end, we have a list of expected utilities that we encountered at each matching. However, this alone may not be that useful, especially if the probability of adverse effect is very low, because these lists will probably look very similar to each other.

So, we do the standard backward induction (from game theory, see, for example Fudenberg and Tirole (1991)). We start from the last matching. We store the average expected utility there and move on to the previous matching. There we take into account the fact that any opponent that does not provide us with at least the average utility of matching 10 is not interesting to us because there is a good chance we might get that utility by waiting until the matching 10.[12] We calculate

[11]This may not always be possible. For example, the opponents may react unfavourably to such information collecting without negotiations. And to get a picture of negotiation tactics employed in the market, it may still be necessary to negotiate.

[12]Of course this is not certain and in fact, in many cases we might be disappointed, but if we are risk neutral (indifferent about a risk) this is how it goes. A more risk-averse approach would be to use the minimum of offers, so that we might be relatively certain of the utility we get in the last matching.

the *average* of the expected utilities in the second-to-last matching, replacing any utilities that are lower than the minimum of the last matching with that minimum (representing the wait until the next matching). This average is then used as the same sort of threshold in the third-to-last matching and the average is again calculated. This average is then used as a threshold in the previous matching and so on until we have these averages for all matchings.

We can observe that the expected utility threshold typically increases when we go towards turn 0 (the first matching). Since the expected utility of the next matching works as a minimum expected utility (because we can always go to that negotiation), it is clear that the expected utility of a matching cannot be worse than any of the matchings following it (because the buyer could always get that later utility by waiting until then). Usually some high quality providers using very buyer-friendly negotiation tactics can provide better utility than the average expected utility of the next matching so the expected utility of earlier matching is often higher than that of later matchings. This is also because at the earlier stages, the buyer still has more opponents to meet and a higher probability of being matched with someone very good. When the matchings go on and no such opponent appears, the probability of that happening dwindles.

Exactly how the expected utility develops over matchings depends on two factors. First, the dwindling possibilities mean (as just explained) that the threshold expected utility is always non-increasing and usually decreasing over time. This means that the decrease in the expected utility threshold is stronger in the cases where the seller can be adversely impacted and the decommitment policy is over-compensatory and decreases slower in other cases. In the extreme cases, where the decommitment fee is very high (*Constant 1.0*) and the probability of adverse impact very high, the threshold may stay unchanged until the very last matching. So what we get is an expected utility for each of the matchings 1-10 that usually decrease over time. These numbers are used to set a *minimum* expected utility that any agreement needs to achieve to be considered. Since the counter tactics usually provide a very accurate estimate for the negotiation's outcome, the useless negotiations can be recognised and withdrawn from immediately. The threshold value in matching m is of course $EU(m+1)$, since if the current negotiation fails, that is the expected utility. After the final matching there is no negotiations and therefore we define $EU(11) = 0$. In the final matching, we therefore have no future offers to be considered.

The future offers are a quite versatile and simple approach that will work basically in any environment and with any of the counter tactics mentioned.[13] So, although it certainly is effective with the *Adaptive Counter* tactic, also other counter tactics can benefit from it. In practise, taking the future offers into account means that when a simple counter tactic has plenty of negotiation partners, it will try to choose ones with a slightly higher quality and ignore the bad ones. And of course the more suitable opponents it will meet, the higher it can set the threshold. Taking the future offers into account may also make the `Negotiator` suggest waiting because higher expected utilities can be achieved later. With the *Adaptive Counter* tactic, the future offers can also mean ignoring opponents that use unfavourable negotiation tactics (such as the *Pure Behavioural* tactic).

Finally, the `Negotiator` handles opponents that fail to meet the expected utility threshold just as opponents with whom no contract can be achieved. It will tell the `Controller` that any negotiation with such an opponent has a zero success probability and if asked to negotiate with such an opponent, it will suggest withdrawing immediately. The `Controller` will usually accept withdrawal in such cases.

9.2 Empirical Evaluation

We will now discuss the empirical evaluation of the `Negotiator`. This discussion is divided into three parts. First we discuss the hypotheses (section 9.2.1) and how the experiments were conducted (section 9.2.2). We then discuss our results (section 9.2.3).

9.2.1 Hypotheses

We start by considering only cases without adverse effects and decommitments, so $a_b = a_s = 0$. We will investigate cases with adverse effects later (see section 10.2.1). We created the counter tactics to be used against the four seller tactics and a fifth tactic, *Adaptive Counter* tactic, combines all four. We therefore contend:

[13] There is no reason to think why a similar approach would not also work with the seller tactics or in fact with any tactic at all. The seller strategies operate on prices and utilities, not expected utilities so in order to make it work in cases with adverse impacts, some additional work would be required.

> **Hypothesis 18.** The four counter tactics work well against their target tactics. The *Adaptive Counter* tactic works well against any of the allowed seller tactics.

This makes the *Adaptive Counter* tactic quite useful in any single negotiation. However, we suspected that the *Adaptive Counter* tactic can sometimes be too adaptive. It will try to reach an agreement with any opponent using any tactic and, by default, it does not consider future negotiations but tries its best to reach an agreement with the current opponent. As explained before, this works well if there is only a limited number of chances for finding a suitable opponent, but if the `Negotiator` gets plenty of chances, it might be better off failing in some negotiations and finding those opponents it can exploit. This can be achieved by using an appropriate counter tactic, since they often fail when the opponent is not using the expected tactic. However, if it is hard for the specialised counter tactics to find suitable opponents, the *Adaptive Counter* tactic will outperform them.

If the buyer succeeds in getting an excellent result in a negotiation only 10% of the time, it is unlikely to do better than an opponent that gets a good or decent result every time. However, if the chance is repeated several times, the situation may well change. For example, a 10% chance in one negotiation changes to 65% if the same 10% chance is repeated ten times.[14] With a 20% chance in each negotiation, the success probability in 10 negotiations is already 89.2%. And this may well start to be useful. Thus we contend:

> **Hypothesis 19.** Other counter tactics can outperform the *Adaptive Counter* tactic if they have a good chance of meeting an opponent that uses the tactic they specialise in.

Even with 10 possibilities, the probability of finding an opponent using the right tactic *and* having a high quality may not be very high. However, the *Adaptive Counter* tactic will be able to get a result with any opponent and if we use future offers to guide it to wait for the good opponents, we should be able to get good results. Thus, we contend:

> **Hypothesis 20.** When future offers are used, the *Adapting Counter* tactic will outperform the competition.

[14]The probability of success at least once in 10 negotiations is: $1 - (1 - 0.1)^{10} = 0.6513$.

However, the peak performance of the *Adaptive Counter* tactic requires that the negotiation tactic the opponents employ is known. When this information simply is not there or is inaccurate, the *Adaptive Counter* tactic will have trouble adapting and, as a result, it will perform less well. This is because the *Adaptive Counter* tactic will try to do well against any opponent and if there is a good chance that the opponent in question might be using one of the behavioural tactics, the *Adaptive Counter* tactic will gear its tactic towards accommodating this and although this means a high probability of success, it will also mean less exploitation and therefore potentially mediocre performance. Because other counter tactics do not care what tactic the opponent is using, but make an offer assuming the opponent uses the relevant tactic, they will not be affected and some of them can outperform the *Adaptive Counter* tactic when information is poor. Thus, we contend:

> **Hypothesis 21.** The performance of *Adaptive Counter* tactic deteriorates as the information about opponent tactics becomes less accurate or is less available. This means that other counter strategies may outperform the *Adaptive Counter* tactic under incomplete information.

9.2.2 Experimental Setup

We run the marketplace 100 times with different settings (different tactic distributions) and with different buyer tactics and add the buyer utilities in different runs together. We then repeat this 100 times and calculate the average utility for the 100 runs instead of one run. This is because in a single negotiation, the variation can be quite large compared to the values. Here the possible outcomes are in the interval $[0.00, 0.50]$ but later when the *Constant 1.00* decommitment policy is used, the interval is $[-1.00, 0.50]$ and the variation will make it very difficult to get any statistically significant results and, therefore, we compare the differences of negotiation tactics over sets of 100 runs. This will give us a result, which says that if we were to enter the market 100 times, then there would be a difference or that there would not be one. There is a good chance that we would be unable to make such a call on a single entry.

Given the average results and variations, we then perform a two-sided t-test to investigate whether or not the differences we can see are statistically significant or not. We will then report the results of these tests where necessary.

In some of our experiments, we use simple tactic distributions where all seller are using the same opponent strategy. This is to show how the different tactics work against a certain tactic. However, in other experiments, we will be interested in how these different tactics work against more heterogeneous seller populations and how the distribution of the seller population tactics affects the performance of different tactics. Because there are four different tactics, changing the relative proportions of each would require significant effort and probably would not produce that much more useful data. So, instead, we have divided the tactics in to two groups of two:

- *non-behavioural*: The *Exponential Time-Dependent* and *Random* tactic and
- *behavioural*: The *Random Behavioural* and *Pure Behavioural* tactic

and we will investigate the effect of the seller tactic composition by changing the proportions these two groups have in the population and divide the group's proportion equally between its constituents. Therefore if say 30% of the population uses non-behavioural tactics (0.3 in the figures), it means that 15% of the population uses *Exponential Time-Dependent* and *Random* tactics each and 35% uses *Random* and *Pure Behavioural* tactics each.

9.2.3 Results

First, we will consider how our counter tactics and especially the *Adaptive Counter* tactic do against different seller tactics. To do this, we ran all nine buyer tactics in a market where the sellers use only one of the four tactics. The results are shown in figure 9.3. As can be seen, against all four seller tactics, the appropriate counter tactic and the *Adaptive Counter* tactic perform equally well and better than any other tactics (this is at the $p < 0.0001$ level in all cases). There is no difference between the appropriate counter tactic and the *Adaptive Counter* tactic in any of the cases, but the wrong counter tactics perform less well. This is especially clear in cases where the sellers are using one of the behavioural tactics and both the *Exponential Time-Dependent* and *Random Counter* tactics fail to produce any contracts (the utility is zero).

The only simple counter tactic that succeeds against any opponent tactic is the *Pure Behavioural Counter* one. This is because the offer that tactic makes (the halfway point between the value and cost) is acceptable to any of the four opponent tactics at some point.[15] Of course the other side of this coin is that it gets quite mediocre results for those contracts it finds. This means that the *Pure Behavioural Counter* tactic gets the same result from all four populations although other tactics' fortunes vary significantly. So it is clear that the simple counter tactics do well against their target tactics and that the *Adaptive Counter* tactic is able to adapt to all four tactics very well when the seller's tactics are known. Both the appropriate simple counter tactic and the *Adaptive Counter* tactic are able to beat all the competition (at the $p < 0.0001$ level). This is consistent with our hypothesis 18 and we can therefore accept it.

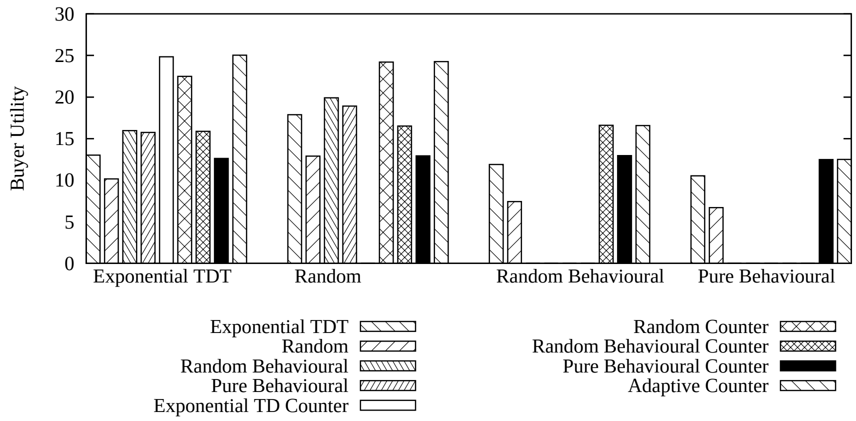

FIGURE 9.3: All negotiation tactics against the single-tactic populations (Hypothesis 18).

Another point worth making in figure 9.3 is that the utility achieved by the best buyer tactic decreases from left to right, so EU(*Exponential Time-Dependent* tactic) > EU(*Random* tactic) > EU(*Random Behavioural* tactic) > EU(*Pure Behavioural* tactic). This holds generally. The behavioural tactics make the sellers using them more difficult for the buyer to exploit. But if and when the buyer can choose between the sellers and knows about the tactics the sellers employ, tactics that are too strong may be counterproductive for the seller. To make any profit, the seller has to get into contracts and by demanding an equal share of the profit (*Pure Behavioural* tactic) the seller may rule itself out of contention, if the buyer has other providers to choose from.

We will now consider the mixed markets, where the sellers use different tactics. As explained above, this means that we vary the fraction of non-behavioural/behavioural

[15]Strictly speaking, this is not entirely certain with the *Random* and *Random Behavioural* tactics, but of course this is very, very likely.

tactics and see how that affects the performance of the different negotiation tactics, the buyer uses. Specifically, figure 9.4 shows how the *Adaptive Counter* tactic performs compared to simple counter tactics. We just showed that the *Adaptive Counter* tactic does at least as well as any other counter tactic when the sellers were homogeneous (that is, when all of them used the same negotiation tactic). However, when the opponents are heterogeneous, the *Adaptive Counter* tactic loses to at least one of them in every single setting we investigated (although it beats all of the simple counter tactics at least twice). This is, as we explained earlier, because the *Adaptive Counter* tactic is very conservative and tries to reach a contract with any opponent. Since it is able to adapt to negotiate with any opponent, it will usually perform very well in any single negotiation. However, the problem is that different opponent tactics yield different types of outcomes. Some are better than others (as was shown in figure 9.3). So basically the *Adaptive Counter* tactic sometimes first encounters an opponent that is using, for example, the *Pure Behavioural* tactic and because of the adaption, it will be able to reach an agreement with that opponent. Another counter tactic might in the same negotiation fail miserably. But the point is that when there is more than one chance of getting a contract, the next opponent might be using a more suitable tactic and this other counter tactic may be able to exploit that next opponent.

The difference between the *Random Counter* and the *Exponential Time-Dependent Counter* tactics comes from two sources. First, the *Random Counter* tactic is able to do well against the *Exponential Time-Dependent* tactic, but the *Exponential Time-Dependent Counter* tactic fails against the *Random* Tactic. This means that the probability of success in a given negotiation is double for the *Random Counter* to that of the *Exponential Time-Dependent Counter* Tactic. On the other hand, when successful, the *Exponential Time-Dependent Counter Tactic* is able to squeeze almost all the benefit from the contract to itself at its opponent's expense. This explains why the *Random Counter* tactic's performance increases faster and reaches a plateau when the fraction of non-behavioural tactics reaches 0.4. This is because at this point, the probability that the `Negotiator` will meet at least one opponent using a non-behavioural tactic is very high, around 99.4%, and further increases in success rate do not bring significant utility improvements. On the other hand, finding an opponent using the *Exponential-Time Dependent* tactic is not yet completely certain (the probability is 89.3%) and therefore the *Random Counter* tactic still has an edge. However, at 0.5 the probability of success for the *Exponential Time-Dependent Counter* reaches already 94.4% and this is compensated for by the higher utility in cases of success so the two perform in

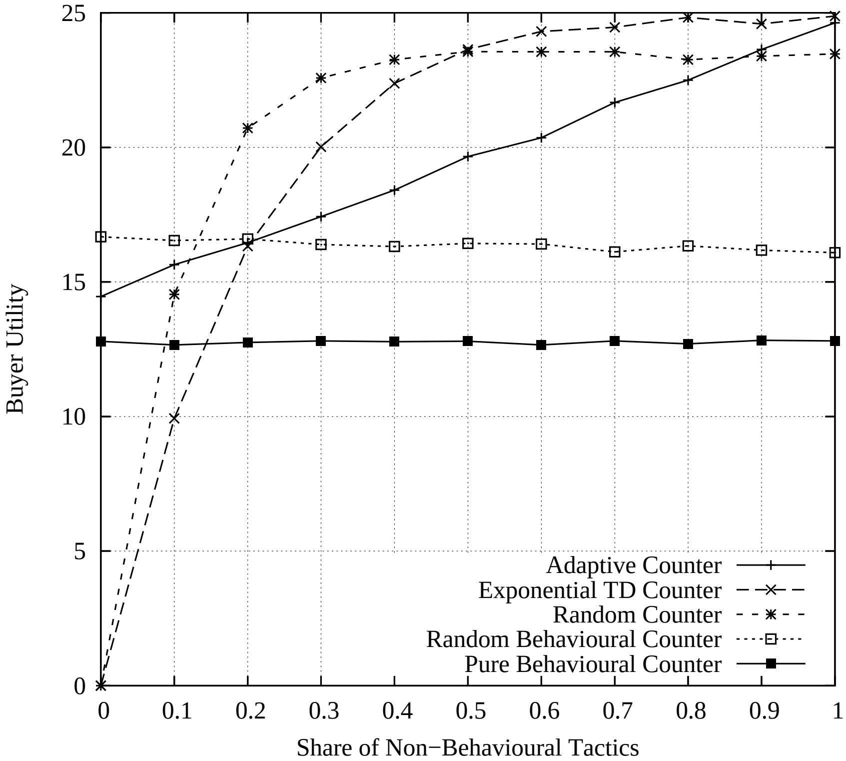

FIGURE 9.4: Counter tactics against various mixed-tactic populations (Hypothesis 19).

similar fashion. So, when the probability increases even further, the *Exponential Time-Dependent Counter* tactic is able to beat all the competition.

The behavioural counter tactics are able to reach a contract against most seller tactics. In the case of the *Random Behavioural Counter* tactic, the only tactic it is hopeless against is the *Pure Behavioural* one, but the probability of finding some other tactic in ten tries is very high in all the cases shown in the graph (the probability is at its lowest when behavioural tactics make up 100% of the opponent's tactic selection, but even then there is a 50% chance that any given opponent uses the *Random Behavioural* tactic, so in 10 negotiations that gives the tactic probability of success of 99.9%). Therefore the seller population composition has very little effect on the performance of these tactics and of course since the *Random Behavioural Counter* actually has a chance of getting more than half of the benefit in any given negotiation, it will outperform the *Pure Behavioural Counter Tactic* in all cases. All these findings are consistent with hypothesis 19, so we can accept it.

However, the above assumed that the buyer agent will take any deal that results in a positive utility for it and this puts the *Adaptive Counter* tactic at somewhat of a disadvantage because it wants to succeed in every negotiation and since in every negotiation, the utility can be positive, it will get itself into contracts that the other counter tactics would not. So, when this tendency is controlled by requiring a certain level of utility for a contract, the *Adaptive Counter* tactic no longer takes any contract from any opponent but only accepts good deals. As before, the *Adaptive Counter* tactic does as well as any counter tactic does in any homogenous population whereas the other counter tactics only do well in one or two of the settings.

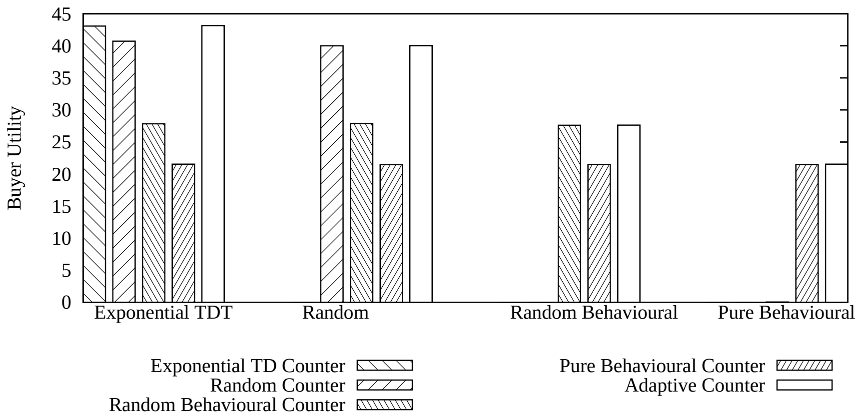

FIGURE 9.5: Counter tactics with future offers against the single-tactic populations.

However, unlike earlier, this advantage translates also to the heterogeneous setting. This is because the *Adaptive Counter* tactic is able to benefit from its adaptiveness and is able to find opponents that are good in terms of both their negotiation tactic and their quality. This is why it is able to beat both behavioural counter tactics, in all cases (at the $p < 0.0001$ level) also when only behavioural tactics are used. The *Random Counter* can find deals but it is unable to exploit them the way *Exponential Time-Dependent Counter* or *Adaptive Counter* can and *Exponential Time-Dependent Counter* may not be able to find high quality opponents that would also use the tactic it was designed for. To this end, hypothesis 20 proposed that the *Adaptive Counter* tactic would be able to beat the competition when future offers are taken into account and this clearly is the case. So we can accept this hypothesis.

Finally, we consider cases with incomplete information. The performance of the different counter strategies are shown in figure 9.7. The specific counter tactics are independent of the tactic information and therefore their performance is a horizontal line. The figure on the left (a) shows how the performance of the

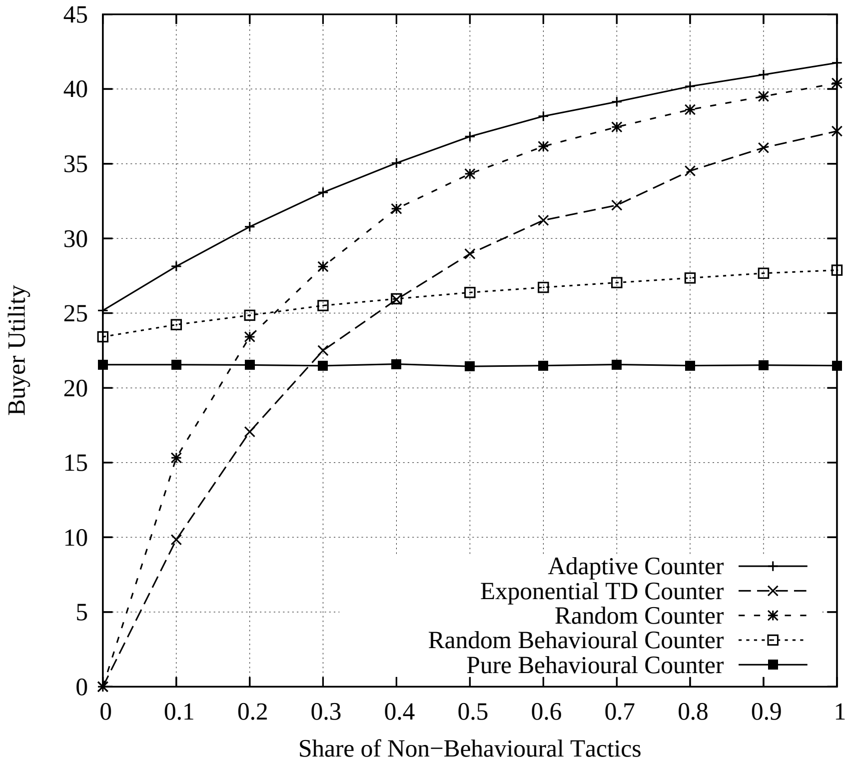

FIGURE 9.6: Counter tactics with future offers against various mixed-tactic populations (Hypothesis 20).

Adaptive Counter tactic degrades as it has less and less information about the opponent tactics. The deterioration comes from two sources. First, as there is less information about the opponents, meeting opponents that the buyer knows the tactic of is gets increasingly less frequent. Second, when the buyer does not know what tactic the opponent uses, the *Adaptive Counter* tactic will often play it safe and make an offer that will be accepted by all opponents. This is because the *Adaptive Counter* optimises the expected utility in a given negotiation and if it makes an offer that is less than what the *Pure Behavioural* tactic would find acceptable, the chance of success drops to 75% (because it will fail with all opponents using that tactic and there is a 25% probability that it is dealing with one right now). Such a drop is usually not compensated for by a 33% increase in the expected utility because the differences between what *Pure Behavioural* and *Random Behavioural* would find acceptable are not large.[16] The best thing to do

[16]Analytically, we can say that the maximum utility for *Random Behavioural* is 2/3rds of the total utility where it is 1/2 for the *Pure Behavioural*. There is exactly the required 33% improvement. However, this the upper bound for the improvement and the actual share that the *Random Behavioural* is able to get is less than 2/3rds. Empirically, we can say that in the pure populations (see figure 9.5), the improvement seems to be around 30%, so very close but not quite enough.

is get the certain deal. This can be seen in the case where there is no information: the performance of the *Pure Behavioural Counter* and *Adaptive Counter* tactics are the same.

Now, with future offers, the buyer will know that there may be opponents with known tactics later and sometimes these might even be something other than *Pure Behavioural* tactics. So the better the chance of that happening, the more likely the buyer will just wait until this happens. If there is no information, the best the buyer can hope for is to have an opponent with good quality instead and, as the results show, this is no different from the *Pure Behavioural Counter* tactic.

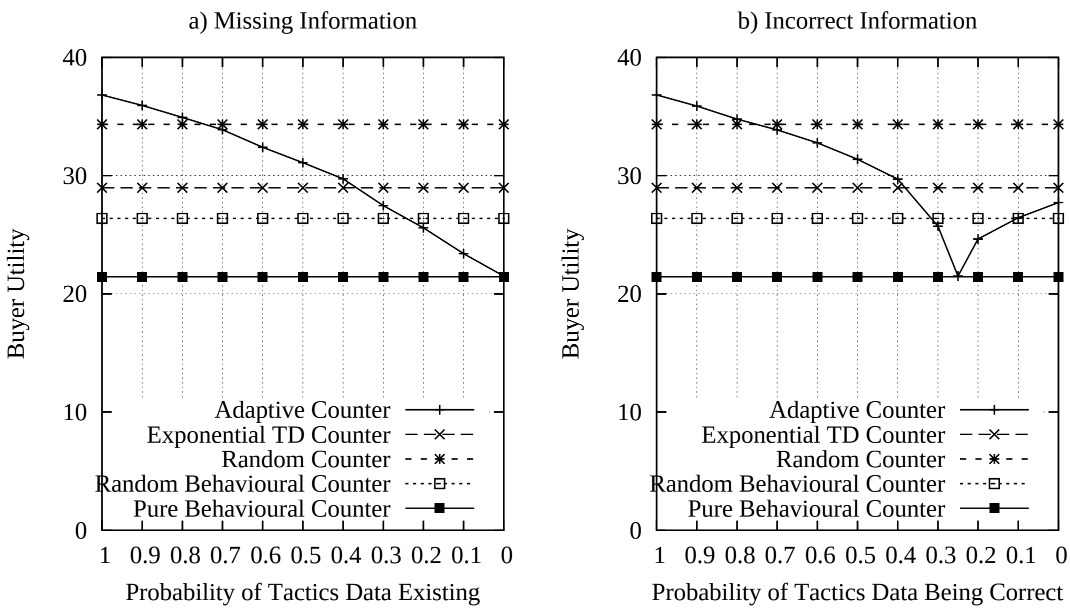

Figure 9.7: Counter tactics with incomplete or incorrect tactic information (Hypothesis 21).

Now, when there is always information, but it is increasingly inaccurate (b), we can observe that there is a minimum at 0.25 and the performance actually improves (quite rapidly) when accuracy gets worse than this.[17] This is mostly because even an incorrect guess gives us more information about the opponent's tactic when we know how accurate our guess is. So when we are accurate at 20% of the time, we know that there is 0.2 chance we are right and the probability that it is one of the three other tactics is 0.8 (or 0.2667 each), so we know that the probability that the opponent is using one of the tactics we did not guess is larger than that our

[17]The exact shape of the curve may not have a sharp point as shown. It is possible, even likely, that there are some values near 0.25 where the performance is also very bad making the point somewhat duller. We did not think the exact shape of the performance curve at this point was that important.

guess is right. This means, for example, that when we guess the opponent is using a *Pure Behavioural* tactic, it sometimes may be enough for us to take our chances and make an offer which will be rejected that 20% of the time. However, if our guess is correct with 25% probability, we could not get any new information about the opponent's tactic, because this probability is exactly what we would get if we took a guess at random and this is what we have when we have no information. When we have some information, we can do better. All these results are consistent with hypothesis 21 and we can therefore accept it.

An important point to note is that the minimum point is not always the same as the performance of the *Pure Behavioural Counter* tactic. If the probability of the opponent using the *Pure Behavioural* tactic is low enough (it seems the critical point is between 20 and 25%), the *Adaptive Counter* tactic would choose to make offers that would be turned down by the opponents using the *Pure Behavioural* tactic. This is why the performance improves very quickly around the minimum point in figure 9.7.b.

We have shown that the appropriate counter tactics can be very effective in a bilateral negotiation and that if there are plenty of opportunities to negotiate, negotiation tactics that fail often but succeed well may be appropriate when the best tactic is chosen. We also showed that an *Adaptive Counter* tactic that adapts to any opponent tactic, needs to consider future offers and have quite a good idea of the opponent tactic to be effective. Finally, when information about opponent tactics is incomplete or incorrect, it can often be better to use one of the simple counter tactics.

9.3 Summary

In this chapter, we introduced the `Negotiator` level of our model and explained a number of negotiation tactics for bilateral negotiation, many of them are novel to at least some degree (contribution **C6**). The new seller tactics are probably the least original, although in exactly that form they have not been discussed in the literature. The simple counter tactics contain some new ideas, like calculating an optimal offer in advance (taking into account the fact that there are many turns in a negotiation) and making that one offer throughout the negotiation to keep the negotiation outcome easy to assess simplifying the problem of giving useful information about negotiations to a problem of assessing success probabilities.

We also introduced future offers or taking into account not only this negotiation but also later negotiations (contribution **C9**), which proved to be very useful even in simple bilateral negotiations. It improved the performance of all counter tactics in many settings, although it of course was not able to make these tactics work where they did not work before. On the other hand, the *Adaptive Counter* tactic was able to benefit most and was able to beat all of the others. The drawback of the *Adaptive Counter* tactic is that it needs relatively accurate information about the opponent tactics to function properly, whereas the simple counter tactics need no such information. An appropriate simple counter tactic should be able to do quite well in any conceivable market, where there are suffificent negotiation opportunities so that an occassional failure is not a problem and where probabilities of success with different offers could be assessed.

We will now move to the next level in our model, the `Controller`, where we discuss how many concurrent negotiations on the same service can be controlled.

Chapter 10

Concurrency Strategies: The `Controller` Level

As explained in chapter 8, the tasks of a `Controller` are:

- *Managing all bilateral negotiations on a single service*: A `Controller` is responsible for starting and ending `Negotiator`s and following their progress and reporting this to the `Coordinator`. A central task is avoiding too many contracts and decommitting from such contracts if they occur.
- *Selecting opponents to negotiate with*: Given the service providers in the market, a `Controller` selects those it wants to negotiate with.
- *Choosing the number of concurrent negotiations*: Given the selected opponents, a `Controller` chooses the number of negotiations to have.
- *Providing estimates on the negotiation outcomes and success probabilities*: To function efficiently, the `Coordinator` may need an outcome estimate of any on-going or future potential negotiations. This estimate requires both the expected utility of successful outcome and the probability of achieving that result and it is a result of combining similar reports from the `Negotiator`s.[1]

In the following, we will discuss how these tasks are achieved. The discussion will relate to our research contributions **C7** (concurrent negotiations and the

[1]This means considering the combined effects of the different negotiations such as multiple contracts (decommitments) and average or worse contracts.

Controller) and, in part, **C9** (future offers). In more detail, we will first discuss the architecture of the Controller (section 10.1). This is followed by discussion of our experiments and their results (section 10.2). We conclude the chapter by summarising our findings (section 10.3).

10.1 Architecture of the Controller

We start the discussion of the Controller's architecture by giving an overview of its structure and interaction with its surroundings (section 10.1.1). We then discuss controller strategies (section 10.1.2). Finally, we discuss how possible offers in later negotiations should be considered (section 10.1.3).

10.1.1 Overview

A Controller controls a group of Negotiators that negotiate on the same type of service with different providers. This means that the Controller decides who to negotiate with and when to accept the offer or quit the negotiation.

The <Controller> component consists of three main parts:

- The <ControllerBase> module deals with communication with the <Coordinator> and the <Negotiator> components. It also starts the new Negotiators and ends the old ones according to its instructions.
- The <Strategy> module's task is to make all the strategic decisions for the <Controller> in a very similar way to the <Tactic> module in the <Negotiator> component, although, as we will see, the <Strategy> consists of multiple strategies, which are all separately interchangable whereas the <Tactic> module was interchangable as a whole.
- The <MarketAnalyser> offers its analysis of the market situation to the <Strategy> component. This module stores the consumer agent's information about the market and the potential negotiations for one particular service.

The interplay between these modules has two different modes. The more complex one, which we call *Opponent Selection* mode, starts when a ControllerBase notices it has no Negotiators but it would still have time to negotiate and it asks the Strategy module if it should start new negotiations (and if yes, who it should negotiate with). This interaction is described in the figure 10.1. Here, after getting a request from the StrategyBase (step 1), the Strategy module will get information about the potential opponents from the MarketAnalyser (steps 2 and 5), who analyses the information it gets from the Marketplace (steps 3 and 4).[2] The Strategy then gives the StrategyBase its recommendation for a negotiation strategy (the opponents and the first offer to make to each of them), including an estimate of the outcome (probability of success and expected utility in case of success) (step 6). The StrategyBase sends this recommendation to the Coordinator for approval (step 7) and the Coordinator responds with the final strategy (opponents this Controller should negotiate with and the first offers it should make) (step 8) which the StrategyBase then implements by starting the required number of Negotiators to manage the negotiations.

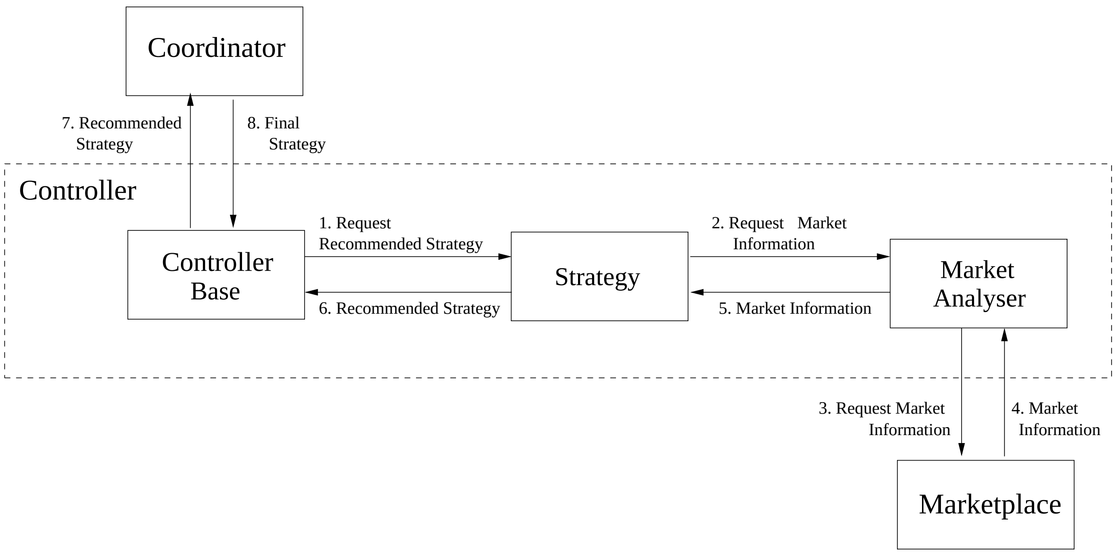

FIGURE 10.1: The *Opponent Selection* mode of a Controller.

The other mode, the *Negotiation* mode, is used during the negotiations (while the Controller has active Negotiators). As shown in figure 10.2, the Negotiators work as a trigger here when they report the progress of their negotiations and suggest what action should be taken next (step 1). When all Negotiators have reported in, the ControllerBase gives these reports to the Strategy and asks it to tell what it should do with the on-going negotiations (step 2). The Strategy responds with a recommended strategy (step 3), which includes an action (accept,

[2]Any call for information from the marketplace is returned only when a new matching is about to begin, because new negotiations can only be started then.

withdraw, counteroffer) to take in each negotiation. This recommendation is then sent to the Coordinator for approval (step 4). The Coordinator returns the final strategy (step 5), which the ControllerBase implements by telling each Negotiator what action to take (step 6).

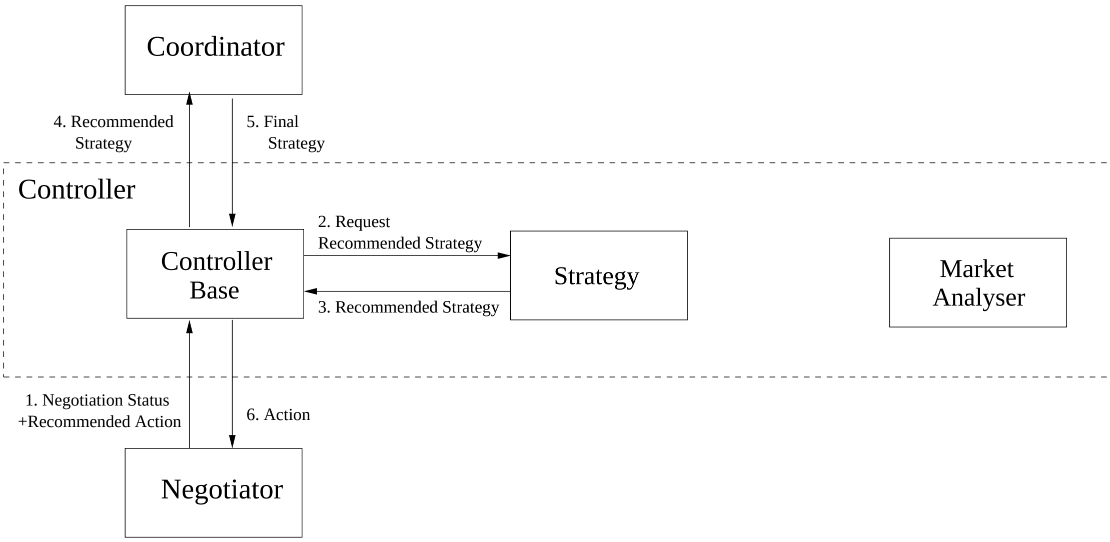

FIGURE 10.2: The *Negotiation* mode of a Controller.

10.1.2 Controller Strategies

In the <Negotiator>, the <Tactic> module was just one component that could be changed to another to get different behaviour. In contrast, the <Strategy> module in the <Controller> is more complicated. It actually consists of four parts:

- The <StrategyWrapper> offers a standard interface for the <ControllerBase> to use and controls the basic interactions between the other components. It is a stable part of the architecture and is therefore not interchangable.[3]
- The <OpponentSelection> ranks the opponents in order of preference as negotiation partners.[4]

[3] Although if this was useful, it would be a relatively straight-forward task to extend the architecture to allow different <StrategyWrapper> modules. This would make it possible to introduce other types of interaction between the other three modules. Currently we can see no need for that, however.

[4] It could also remove useless opponents and although we implemented this functionality (Opponent Filtering), we do not use this functionality in the experiments we are going to discuss. Therefore it will not be mentioned again.

- The `<ConcurrencyControl>` decides how many negotiations there will be concurrently.
- The `<NegotiationStrategy>` then decides what to do in each on-going or new negotiation.

All modules except `<StrategyWrapper>` offer a standard interface and can be changed, although the `<OpponentSelection>` and `<NegotiationStrategy>` modules actually consist of more than one interchangable part. This structure allows very flexible selection of strategies and also makes it very easy to add new strategies.

As explained, the `Controller` has two modes of operation: Opponent Selection and Negotiation. In the former, all modules are used as shown in figure 10.3. Basically, the `ControllerBase` asks the `StrategyWrapper` to provide a status report and a recommendation for the strategy. `StrategyWrapper` calls in turn:

a. `OpponentSelection` (step 2): It gets the market information from the `MarketAnalyser` (steps 3 and 4). Among other things, this information contains the list of potential opponents. The `OpponentSelection` sorts the opponents into order of preference and may remove some opponents from the list altogether. The sorted list is returned to the `StrategyWrapper` (step 5).

b. `ConcurrencyControl` (step 6). It decides how many negotiations to run concurrently. The list of negotiations (with opponents) is returned to `StrategyWrapper` (step 7).

c. `NegotiationStrategy` (step 8): It decides what to do in each of the ongoing negotiations, calculates the expected result of these negotiations and sends this information to the `StrategyWrapper` (step 9), who forwards it to the `ControllerBase` (step 10).

In the *Negotiation* mode, the interaction is much simpler and consists only of the `NegotiationStrategy` module checking on the progress of the negotiations. This includes handling any contracts, doing any decommitments necessary and deciding whether or not to accept an offer, withdraw from a negotiation or what offer to make.[5] The recommended strategy here consists of the next actions to take in each negotiation, expected outcome (success probability and utility) and

[5]The information flows through the `StrategyWrapper` module but it only calls the appropriate method in `NegotiationStrategy` and returns the outcome to the `ControllerBase`.

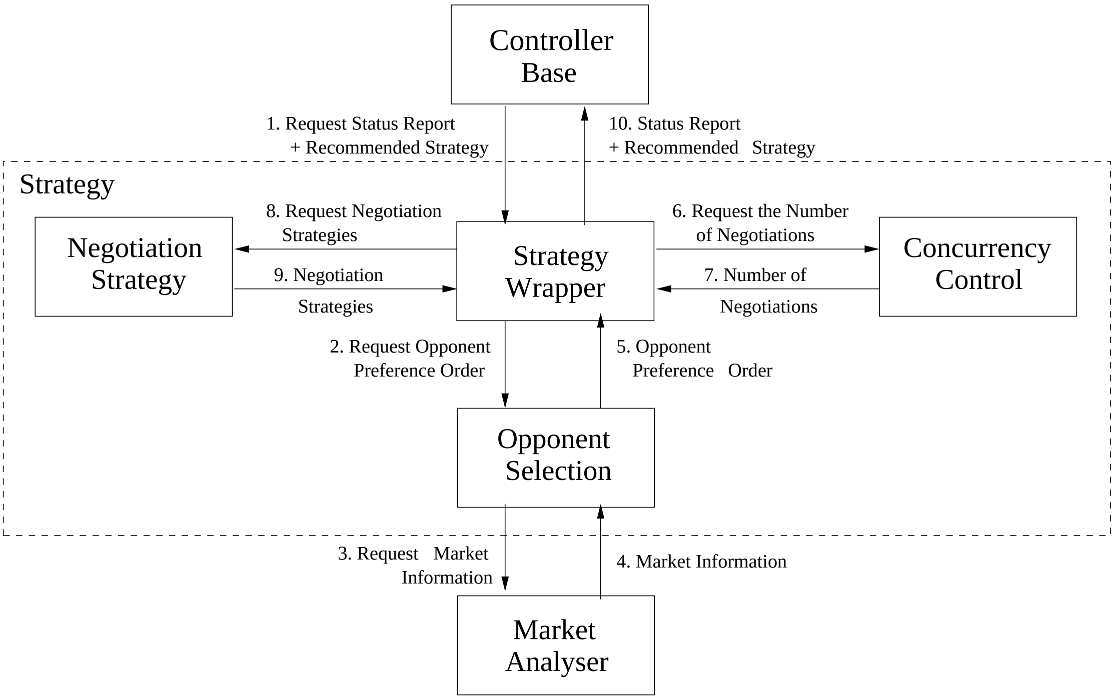

FIGURE 10.3: The `<Strategy>` module in the *Opponent Selection* mode.

the list of decommitments to make. In both modes, the `ControllerBase` sends this recommendation to the `Coordinator` to be checked against other `Controllers`' strategies.

We will then discuss the different strategies in more detail. First, we will discuss opponent selection strategies (section 10.1.2.1). Second, we will discuss methods for deciding the number of concurrent negotiations to have, that is, the concurrency strategies (section 10.1.2.2). Finally, we discuss negotiation strategies (section 10.1.2.3), which set the negotiation tactics in each individual `Negotiator`.

10.1.2.1 Opponent Selection Strategies

As discussed in the previous subsection, the `<OpponentSelection>` module does not actually do any selecting, but only arranges the potential opponents according to some rule in decreasing order of relevance, placing the best opponents in front and the worst towards the end. We have three basic opponent ranking rules:[6]

- *Random* opponent selection means that the order of opponents is chosen among the sellers in the market at random. The ordering of the opponents

[6]The different strategies are implemented by separate modules that provide a standard interface for the `StrategyWrapper` to use.

is re-done at each matching so that all sellers have an equal chance of getting to the top no matter when they entered the market.

- *Quality* opponent selection means that the sellers are ordered according to the quality of the service they provide, in descending order (the providers with highest qualities first). Because the expected utility usually increases with quality, this usually means improvement over the *Random* opponent selection. However, because the different negotiation tactics that different opponents may use can affect the expected outcome quite significantly, there is still some variation in the expected utilities.

- *Expected Utility* opponent selection means that the opponents are ordered based on the expected utility they would offer in a negotiation. This means that the opponent selection asks the `Negotiator` level to provide an estimate for a negotiation outcome on each opponent and then uses these estimates to order the opponents in descending order, the opponents offering the highest utility first. This means that the first few opponents are usually very high quality providers that are using a negotiation tactic that is very suitable to the negotiation tactic the `Negotiator` uses (if such information is available).[7]

In the *Expected Utility* opponent selection strategy, calculating an expected outcome of an opponent means devising an optimal tactic against it and estimating the success probability of this tactic. Obviously this can only be done by the `Negotiator` who is expert in bilateral negotiation.

On the other hand, it is equally clear that calculating the optimal tactic is not a trivial operation and that we do not want to do it if we do not have to. Doing it for several dozens of opponents at each of the ten matchings[8], we would be doing much of the work for no good reason since we are unlikely to negotiate with most of the opponents. Therefore, we use optimisation to minimise the number of opponents we have to estimate the expected utility with. This is done by calculating (offline) an upperbound for the expected utility given the quality the opponent provides, the negotiation tactic (if known), the adverse effect probabilities and the decommitment policy. If the upperbound of an opponent's expected utility is lower than the threshold utility, it will be ignored (its actual expected utility is not calculated). The threshold is the nth best expected utility we have found so far, where n is the maximum number of opponents we may want to negotiate

[7] If tactic information is not available, only the quality is used.

[8] The expected utility will usually be different at each matching even if the opponent is the same, because the effect probabilities will decrease over time.

with. So if we know the opponent is not going to make it to the top n, there is no reason to calculate the exact expected utility for it and we can safely ignore it. And to help find a good threshold quickly, the opponents are first ordered by their quality, so higher quality providers will be investigated first. If and when our estimates are reasonably accurate, this can cut down the number of opponents to be processed significantly.

10.1.2.2 Concurrency Strategies

The `<ConcurrencyControl>` module selects the number of bilateral negotiations to use. It is given the ordered list of opponents and it then proceeds to decide who to negotiate with. It cannot pick and choose the opponents from the list, but if it wants to negotiate with the fifth opponent on the list, it will have to negotiate with opponents $1 - 4$ too. Giving a concurrency strategy the right to pick opponents from the list would mean stepping into opponent selection, which is not the domain of the concurrency strategy. It will take the list as given and just choose the number of opponents from the top of the list. We will now introduce the three families of concurrency strategies we have.

A *Simple* concurrency strategy always starts a fixed number of negotiations at each matching. We use ten variations of this group, namely from *Simple 1*, *Simple 2*, ..., *Simple 10*, where the number specifies the number of negotiations it starts at each matching, so how many concurrent negotiations we will have. The *Simple* concurrency strategies do not take a risk of multiple contracts or success probabilities into account in any way. Rather they always start the specified number of negotiations as long as there is that many opponents in the market. If there are fewer opponents in the market, they start negotiations with all the remaining ones.

The second type of strategy is the *Analytic* concurrency strategy. This represents an improvement over the *Simple* strategies because it will explicitly take into account the risk of getting mediocre or multiple contracts and success probabilities. As the name suggests, the *Analytic* concurrency strategy calculates the expected utilities for each and every negotiation number allowed and it chooses the one that maximises the expected utility. How this happens in detail is complicated because with multiple negotiations, there are many different cases to consider.

To demonstrate this, let us take an example of three negotiations. Here, the basic possibilities are that we may have $0 - 3$ contracts at the end of the negotiations.

We have to consider each case. If there is no contracts, we have a zero utility, so that case is simple. If we have one contract, it could be from any of the three negotiations. Thus, we have to calculate the probability for each negotiation (and that negotiation alone) producing a result and then have an expected utility for each case. This is simple and the information (the chance of success and the expected utility in case of success) is provided by the `Negotiator`s.

After that, however, things become complicated. If we have two contracts, we can of course calculate the probabilities of that happening for each case, but it may well be that one contract is formed clearly before the other and we can stop the negotiations before we get the other contract (we would have to decommit from them). To do this, we have to estimate the probability that we will actually have two contracts. This is done by estimating a contract time distribution for each negotiation. The different negotiation tactics and their countertactics have different patterns when it comes to contract times (if we use the *Adaptive Counter* tactic):

- *Exponential Time-Dependent*: the contract is usually formed at the seller's deadline, when the seller would offer its reservation price that the buyer has been offering all along. However, if the seller uses a very strongly conceding variation, it might actually reach the reservation price earlier (because we have rounded our offers). Also, the seller might not make the reservation price offer at the last turn of the negotiation, but it may also have a deadline during the negotiation and if that is the case, it will offer the reservation price then. The probability of the seller having a deadline during the negotiation can be estimated and very conceding negotiation tactics can usually be ignored (see below).
- *Pure Behavioural*: Here the countertactic will have the buyer making an offer that the seller cannot refuse on the first turn and therefore the contract is usually formed there and then.
- *Random Behavioural* and *Random*: Here, the countertactic will make an offer that has a relatively small chance of getting accepted at one turn, but when made up to 100 times (maximum length of a negotiation[9]), the probability that it will be accepted at least once is often very high. Ex ante the probability is at its highest on the first turn and decreases slowly towards the end (because there is an increasing chance that a contract has been found earlier).

[9]See section 8.1.1

With the *Random Counter* tactic, we can only succeed against *Exponential Time-Dependent* and *Random* tactics. With the *Random* tactic, the situation is identical to when we use the *Adaptive Counter* tactic. With the *Exponential Time-Dependent* tactic, it depends on the β parameter value but there is a peak towards the end of the distribution (for simplicity we assume that the contract will be formed at the last turn). In all cases, we also have to take into account the deadline distributions of the opponents. If we are in the first negotiation, there is a 90% chance that the opponent's deadline is later than turn 100, meaning that its deadline in this negotiation will be 99. But if we are in the last negotiation, there is only a 1% of this happening.

Armed with these distributions, we use them to actually calculate the probability of having the two negotiations end in a contract at the same turn. This is of course:

$$\sum_{t=0}^{100} P(\text{contract in negotiation 1 at turn t})P(\text{contract in negotiation 2 at turn t})$$

This gives us the probability that we have to decommit from one of the contracts. We then have to do a similar calculation for the other two cases (negotiations 1 and 3 and negotiations 2 and 3). In each case, we consider the fact that in the eventuality that we have two contracts, we can choose the better of the two, but that we would have to decommit from the other one. Again, we get probabilities and expected utilities.

The case where we have three contracts is similar and again we can calculate the probabilities of this happening at the same turn in all three:

$$\sum_{t=0}^{100}(\prod_{n=1}^{3} P(\text{contract in negotiation } n \text{ at turn t})).$$

To get the expected utility for this case, we take the expected utility for the best contract and deduct the decommitment fees for the other two.

To calculate the total utility for the three negotiation case, we then add all these expected utilities (probability times the (expected) utility) together. We proceed to do this for all possibilities from 1 to whatever maximum we have assigned and once we have estimates for all cases, we choose the number of negotiations based on what we estimate will give us the highest expected utility.

This is a reasonably computation-intensive process and the worst part of it is that the number of cases will quickly increase with the number of negotiations. With one negotiation, we only have two cases (success and failure). With two negotiations, we have four cases, both negotiations successful at the same turn, only negotiation 1, only negotiation 2, and neither successful. With three, it's 8 (=1+3+3+1) cases, four 16 (=1+4+6+4+1)[10] and so on. At some point, we will run out of computing power and we have to discover some heuristics. These could have to do with known properties of negotiation tactics[11] Or we could calculate such probabilities for any mixture of opponent strategies in advance, so we would know the probabilities of the same round successes if we have 3 opponents using the *Pure Behavioural*, 4 sellers using the *Random* and 1 seller using the *Exponential Time-Dependent* tactic when we are using *Adaptive Counter* tactic on each of them. Here, we use a relative small maximum number of negotiations (10 which gives us 1024 cases on that level, a large but still manageable number) and the calculation, although somewhat involved, can still be managed.

So, the *Analytic* concurrency strategy adds together all the cases for up to the maximum number of negotiations (as we just mentioned, this will be 10 in our experiments), calculates the expected utility for each number of negotiations and then chooses the one that provides the highest expected utility. As with the *Simple* strategies case, the *Analytic* concurrency strategy does not change the negotiation tactics that the `Negotiators` come up with. Rather it takes the `Negotiators`' tactics as given and tries to optimise the number of negotiations to have.

The good part of this approach is that it can also be used under incomplete information. If there is, for example, a 50 : 50 chance that the opponent is using a *Random* or *Pure Behavioural* tactic, once we know our own tactic, we can guess a time distribution for any outcome and then use this in our calculations. Also small errors here and there are not likely to significantly affect the outcome, because they only mean that we might be slightly over- or under-eager to have more negotiations

[10]The numbers are of course binomial coefficients, which are often described by Pascal's triangle. This is because the problem here is just a type of binomial distibution (each negotiation either fails or succeeds just like you get heads or tails in a coin toss).

[11]For example, any analytical approach with *Adaptive Counter* needs to consider at least three cases: two or more opponents using the *Exponential Time-Dependent* or the *Pure Behavioural* tactic will lead to two or more contracts and should usually be avoided (if the fee is positive). However, because the contract with the the *Exponential Time-Dependent* tactic will be usually formed very late, having opponents using different tactics can usually decrease the problem because they make it likely that a contract will be formed earlier in these other negotiations. Moreover, the probability that the *Random Behavioural* or the *Random* tactics form a contract at the same turn is usually quite small, although should not be ignored especially if there are more than two concurrent negotiations.

and although this may have an effect on our performance, we should usually still do better than doing something simple like having a fixed number of negotiations. Although of course there may be cases where a single number of negotiations is the best approach in all cases.

The *Adaptive* concurrency strategy does not take the `Negotiators`' suggestions as given, but it will choose the best level of offers in each negotiation separately, together with the `Negotiators`. This will allow the *Adaptive* concurrency strategy to avoid getting into inferior contracts, which can be a problem for other strategies, when the expected utility varies strongly among the opponents the buyer encounters. Basically the *Adaptive* concurrency strategy expects the same expected utility from each successful negotiation and sets its offer in each negotiation so that all negotiations provide the same expected utility in case of success. The *Adaptive* strategy will calculate the offer needed in each negotiation and then ask the `Negotiator` level to provide the success probability with that offer. These probabilities and expected utilities will then be used in an identical process to the *Analytic* concurrency strategy to set the number of opponents optimally. The *Adaptive* concurrency strategy therefore varies not only the number of negotiations, but also the expected utility from each negotiation and that makes it possible to avoid some of the pitfalls the other strategies are not able to avoid.

The *Adaptive* concurrency strategy may try different levels of expected utility before it is happy with the results. It starts from the best expected utility of a single negotiation (any of the 10), and chooses the optimal number of negotiations to have given the success probabilities in the negotiations the `Negotiator` has provided. It then tries a bit lower and higher expected utilities to see if it could increase the expected utility that way. It therefore finds a local expected utility maximum (which is also likely to be the global maximum) and uses that.

The problem with the *Adaptive* concurrency strategy is that it assumes complete information about opponent tactics and given this, it does not function properly, if the actual expected utility of the best offer is not known. This is because it starts looking for the optimal expected utility to demand in each negotiation at the best opponent's level and finds the local maxima nearby. However, with incomplete information such local maxima may not be a global maxima and, therefore, the best offer it finds may be far from optimal. For example, the *Adaptive Counter* tactic would under incomplete information often suggest making an offer that all opponents accept and the highest expected outcome would be offered simply by the opponent offering the highest quality. The problem is, unlike the case with full information, most of the other opponents would be able to provide that sort expected utility with the varying level of probability and this means it will be much more difficult for the *Adaptive* concurrency strategy to find the optimal offer. There are several problems with using the *Adaptive* concurrency strategy under incomplete information:

- There is always a chance that more than one of the opponents will be able to match the required expected utility level even if they actually did not. This means the *Adaptive* concurrency strategy might demand too high expected utilities (risk missing contracts). On the other hand, it may also underestimate the probabilities of extra contracts and get into more contracts than would be useful (risk getting too many contracts). This is the direct result of incomplete information.

- More critically, however, it may well be that the starting point is much further from the global optimum than it was with the complete information and, on the other hand, the combined effect of several negotiations is much more difficult to estimate with incomplete information. This may mean that there are local maximums before the global maximum. If the *Adaptive* strategy gets stuck in a local maximum that is much lower than the global maximum, it may well mean that most opponents are still in contention and especially with high decommitment fees, the *Adaptive* strategy may be forced to play it safe and have very few or even just one negotiation. This one negotiation may not even be the best one available, but a very mediocre one.

In other words, the *Adaptive* concurrency strategy, as we have implemented it, loses much of its advantage (ability to pick out the good opponents from the mediocre, under incomplete information such division is much more vague) under incomplete information. The fact of the matter is that in the few experiments we conducted the *Adaptive* concurrency strategy as we had it, does not seem to be particularly impressive under incomplete information. This may be in part an implementation issue but it will also be in part, because the approach we took works markedly less well under incomplete information. We have decided to leave adapting the *Adaptive* concurrency strategy to use incomplete information to future work (see section 12.2.2 for more detail) and in this work, we will use the *Adaptive* concurrency strategy only in the settings where full tactic information is available.

10.1.2.3 Negotiation Strategy

The `<NegotiationStrategy>` module approves the decisions on all bilateral negotiations, it decides when offers are accepted, when to withdraw from a negotiation and what offers to make. It also manages contracts that have been formed and handles decommitments, if necessary. This is the only module that is active in both modes of operation: As explained in section 10.1.2, it will give the first offers to make in new negotiations in the *Opponent Selection* mode and will suggest plans of action for each negotiation in the *Negotiation* mode.

There are clearly separate areas of this negotiation strategy and, therefore, the `<NegotiationStrategy>` module is further divided into three parts:

- The `<AcceptStrategy>` makes all decisions related to contracts. So, it makes the decision on whether or not to accept any of the offers the opponents made this turn (and if yes, which one to accept). In addition, it will also decide what to do with contracts that have been entered into, because one or more of the opponents accepted the buyer's last offer.

- `<WithdrawalStrategy>` decides whether or not any of the existing negotiations should be withdrawn from.

- `<CounterofferStrategy>` decides the counteroffer to make in each negotiation.

The interplay of these three modules is simple. First, the `AcceptStrategy` checks if any contract has been formed or if there are any offers worth accepting. Second, the `WithdrawalStrategy` checks if there are any negotiations that should be withdrawn from. Finally, the `CounterofferStrategy` decides what counteroffers to make in the remaining negotiations. As earlier, each of these submodules provides a standard interface to allow new strategies to be implemented.

In this work, we only have one of each strategies and they are all somewhat basic. The *Simple* accept strategy first checks if there are any new contracts (the buyer's offers that a seller has accepted). If there is more than one contract, the best one is chosen and others are marked as contracts to be decommitted from. If there is none, it will move to consider if any of the offers made by the sellers is good enough to accept, i.e. give an expected utility above the specified threshold (zero or whatever future offers say). If there is more than one such offer, only the one giving the highest expected utility is chosen.

The *Simple* withdrawal strategy will see if the buyer has found a contract. If it has, it withdraws from all remaining negotiations. This also happens when a deadline is reached or if a negotiation cannot lead to a good result (`Negotiator` says such a result cannot be obtained).

The *Simple* counteroffer strategy will take its lead from the recommendation made by the `Negotiator` or the *Adaptive* concurrency strategy if such a recommendation exists (on-going negotiation). If it is the first round and no recommendation is forthcoming, it will call the `Negotiator` level to ask for the optimal offer and use that instead. The *Simple* counteroffer strategy will not start second-guessing recommendations made by others who are more qualified to make this decision. Its task is only to make sure that a counteroffer for each continuing negotiation is always available in the recommended strategy.

10.1.3 Considering Future Offers

It seems likely that we could benefit from taking the future offers into account also when we allow multiple concurrent negotiations.[12] Of course, the fact that we have different opponent selection and concurrency strategies means that we have to adapt our approach to take these into account. Fortunately, this is reasonably straight-forward and the basic approach is exactly as before. We run the market but instead of negotiating, we collect the information about the opponents that the buyer would see if it would enter into such market. In other words, we calculate the expected utilities for each and every opponent we might negotiate with and store this information. We then use backward induction to determine the appropriate threshold levels just as before (see section 9.1.3).

However, the different concurrency strategies make it a slightly more complicated process. We have twelve different concurrency strategies (*Simple 1-10* and *Analytic* and *Adaptive*), but only eleven different ones in terms of what opponents they might see, because the advanced strategies (*Analytic* and *Adaptive*) are very similar in this respect. Fortunately, we can get the values for all these eleven cases during the same run, so we only need to gather data once in each setting (opponent selection/effect/decommitment policy setting).

As already mentioned, we collect the information by employing a specific concurrency strategy that never starts any negotiations but just collects information

[12] We did this with one negotiation at the `Negotiator` level in section 9.1.3.

about the opponents. It starts each matching by taking the list of opponents it gets from the opponent selection and it goes through the list (no opponent is ever removed during the process) opponent by opponent, calculating the expected utility it offers and then going through the concurrency strategies one by one checking if that particular strategy has already negotiated with this opponent and, if not, stores the expected utility of this opponent and marks it as an opponent this strategy has already negotiated with, so it will not be considered again with that strategy in the later matchings. Obviously in the *Simple 1* strategy, only one 'new' opponent can be considered at each matching, whereas in the *Simple 10* strategy, ten new opponents need to be found and in the later phases, a large number of opponents may have already been negotiated with. The only exception to this process is the *Analytic/Adaptive* process which always considers the first ten negotiations because the best approximation on the behaviour of these strategies is that they do not negotiate until they find a good deal and then they exit with that deal. When all strategies have found their maximum number of negotiation partners, the process moves on to the next matching.

We then come to the second essential difference to the approach we had in the bilateral negotiation. Here, we store separately up to 10 different values for each matching of each run. We take the best utility found, the second-best utility and so on up until the number of negotiations we may have (10 for the *Analytic/Adaptive* and n for the *Simple n* strategy). This is because sometimes it might not be the best approach to only have a few opponents due to a very high threshold, but it may sometimes be a good idea to have a slightly lower threshold (and more opponents). We could have of course reached a similar approach (trying lower thresholds too) by using something other than average as our means of deriving the threshold levels or even simply by trying say 5% lower thresholds and see how that works. However, this was the most straight-forward way of lowering the thresholds and still using the data collected in a useful way. Also, it seemed the best approach. A five percent drop in thresholds might be a huge amount in some settings and nothing in others. A similar problem is connected with using some other method than average in the calculations. Considering only the best options may not give an accurate picture of the actual opponents the buyer meets. The difference between the best and the second or fifth best might in some cases, be huge and, in others, very small.

One particular problem associated with the cases with no tactics information is that the thresholds will be slightly different when they are collected with and without the tactics data. For example, consider the *Random Counter* tactic. It will succeed only against sellers that use non-behavioural tactics, which in our setting are around 50% of the population. Now, in a setting where the qualities vary significantly (the *Random* opponent selection), this means that without tactic information, all opponents are considered when the thresholds are calculated (because it is unknown which opponents are using which tactic), but without tactic information, only the opponents using the non-behavioural tactic will be considered. That is because some of the high-quality opponents we encounter while collecting data are using one of the behavioural tactics and therefore would be ignored when we have the tactic information but they are not ignored when we do not have the information. Because without the tactic information, the expected utility of an opponent depends only on its quality, this effectively means that without the tactic information, the quality threshold is slightly higher, the buyer expects a higher quality before it negotiates.

This effect can be beneficial in some cases and can even mean that tactics such as *Random Counter* will do better without the tactic information than with it. This is because it may mean that more negotiations are used under incomplete information and, therefore, more options will be investigated. The highest quality threshold is unlikely to be hugely successful because it will mean very few opponents and some of them will not be using a suitable tactic. However, when the threshold is decreased from this top level, it may well be that there are many more opponents in the market that get over the buyer's threshold and even if negotiations with many of them will not be successful (because they use behavioural tactics), it will mean that the buyer encounters more opponents in the end and that can be useful in some circumstances. We will discuss this further in section 10.2.1.3.

In the case of the *Adaptive Counter* tactic, this effect may be different. The nature of the effect depends on what type of offers the tactic chooses to make. If it makes offers that any opponent can accept, then there is no similar effect as with the *Random Counter* tactic, but of course, if it should make an offer that opponents using some of the tactics reject, the situation would be analogous to what we just discussed.

To sum up, in this chapter, whenever we say we use 'future offers', it means that with all strategies we have used all different future offer threshold levels described here and only consider the one that produced the highest expected utility in our experiments. Of course our approach does not mean that we always find the threshold levels that produce the highest possible outcome, but we believe our approach finds reasonable levels and by trying up to ten different possibilities, we are likely to get a better idea of the performance of different strategies than if we just used one. Since we use the same approach in all cases and strategies, it seems reasonable to say that any differences we might see are real and not just coincidences. We will not usually report which threshold levels worked best in each and every experiment, because that would just take a lot of space and is generally speaking not that interesting. However, we can say that with low effect probabilities, the first-best thresholds usually yield the best results. But when the buyer effect probability is medium to high, also the lower thresholds can be useful. We do not make any detailed proclamations on how these thresholds should be set in a given environment. We have simply described one method and use it to improve our results. However, we can say that sometimes it makes sense to experiment with several levels of thresholds to see what works and what does not.

10.2 Empirical Evaluation

We have discussed the `Controller` and its parts and some strategies it can employ. We will now investigate how the controller and these strategies work in practice, what impact they will have on the buyer's performance. As in the earlier parts, this section is divided into three parts. First, we discuss theory and make our hypotheses (section 10.2.1). Then we will discuss how we are going to investigate these hypotheses in detail (section 10.2.2) and finally we provide our results (section 10.2.3).

10.2.1 Hypotheses

Here, we will be most interested in two main things: (i) concurrency strategies (especially the advanced strategies versus the *Simple* strategies, but also differences between the two advanced (*Analytic* and *Adaptive*) strategies), and (ii) opponent selection. In general, we expect the *Analytic* and especially the *Adaptive* strategies to do well because they take into consideration more factors and we expect more sophisticated opponent selection strategies to help the buyer to do better. However, in our setting, the concurrent bilateral negotiation with adverse effects, is so complex that we do not expect these general ideas to hold in all circumstances. It seems likely that sometimes the advanced strategies may not be that helpful and there might even be some circumstances under which they are even counterproductive.

Before we can discuss any detailed hypotheses, we need to explain the range of situations we are going to explore here. There is going to be many of them and it will be somewhat difficult to make very detailed hypotheses that would be relatively compact but still describe the different situations we face. On the other hand, we need to explore many settings to see where our approaches work and where they might not work.

So, we use two different decommitment policies, one undercompensatory (*Constant 0.0*) and one overcompensatory (*Constant 1.0*). These two were chosen to test our strategies in a setting with very different decommitment policies and for their simplicity. We have two different negotiation tactics (the *Random Counter* tactic and *Adaptive Counter* tactic). These two were chosen because they were the two best negotiation tactics in the bilateral negotiation in the setting we use (all seller tactics equally likely).[13] Moreover, they offer different types of approach and it will be interesting to see what differences, if any, that causes in the more complicated setting.

In addition, we have two different information settings. In the first, we have full information about the negotiation tactics the sellers use and, in the second, we have only the tactic probability distribution in the seller population. And we have settings where the buyer agent has varying estimates for the future offers to expect and where it has none. And of course in every one of these we then have 28 settings, where buyer only, seller only or both of them can experience an adverse effect with a probability of $0.0 - 0.9$.[14]

Each of these settings and combinations of these settings will have its own characteristics. For example, when the buyer is potentially affected, a higher decommitment fee makes it necessary for the buyer agent to be more careful about entering into contracts because of the high cost of getting out of commitments, but when only the seller can be affected, the non-performance is actually preferrable to the buyer (due to the over-compensation) and the buyer may want to ensure getting a contract as soon as possible. On the other hand, the *Adaptive Counter* tactic with full tactic information succeeds basically with any opponent whereas the *Random Counter* tactics succeeds only with opponents using the *Exponential Time-Dependent* or *Random* negotiation tactics and fails to reach an agreement

[13] These results we discussed in section 9.2.3.

[14] There are nine cases ($0.1 - 0.9$) for each of the three situations and one common case for all, where the probability of adverse effect is zero for both parties. The cases with certain effects (1.0) are not that interesting and so are not considered here.

with everybody else.[15] This obviously has implications on the number of negotiations we should have. The problem with high number of negotiations is, of course, getting into too many contracts and having to decommit from the extra contracts. This can be very expensive when the fee is very high. On the other hand, if decommitting is free, then the buyer does not need to worry about extra contracts all that much. Also, we use the information about the future offers as a utility threshold for negotiation. Therefore, if a `Controller` knows that it will probably be able to find a better contract later, it will either not negotiate with bad opponents at all or instruct the `Negotiator` to fail in those negotiations. This will make the difference between the *Adaptive Counter* and the *Random Counter* tactics smaller, although it does not remove it all together. Having full information about opponents' tactics and having only the probability distribution makes the situation quite different for the advanced concurrency strategies. Their ability to predict extra contracts and successes drop significantly and this will probably show in their performance.

And of course the two or three opponent selection strategies (*Random*, *Quality* and *Expected Utility*) we have, will have their own impact on the setting. The *Random* opponent selection means that the buyer encounters opponents that have a large variety in both quality and negotiation tactics, whereas *Quality* opponent selection means meeting only high quality opponents and most of the variation is in the negotiation tactics. The *Expected Utility* opponent selection removes much of the negotiation tactic variation too, since with the two negotiation tactics we have, it means meeting mostly high quality opponents using either the *Exponential Time-Dependent* or *Random* negotiation tactics, because such opponents typically provide the highest expected utility. This will obviously have a strong impact on the performance of different concurrency strategies. Sometimes it makes sense to have many negotiations, sometimes only a few. Sometimes it is a good idea to wait for the probability of adverse effect to go down before starting to negotiate, sometimes it does not matter or might even be harmful, and so on. We have 11-12 concurrency strategies in each setting (10 different *Simple* strategies plus the *Analytic* and *Advanced* strategies).

As mentioned, we have two different information settings: one with full tactic information and one with no tactic information. In addition, we may or may not use future offers. We, therefore, would get four different settings. However, we will only discuss the case with future offers when we have no tactic information,

[15]Moreover, the *Adaptive Counter* tactic is able to exploit the *Exponential Time-Dependent* tactic slightly better than the *Random Counter* tactic.

because the setting with no tactic information and no future offers would not bring anything that new to the table. The effect of future offers is similar (although a bit weaker because of less information) both with the tactic information and without it and, therefore, there is little reason to repeat it. This means we will discuss three settings. First, we discuss cases with full information about opponent tactics but no future offers (section 10.2.1.1), then full tactic information with future offers (section 10.2.1.2) and, finally, no tactics information but with future offers (section 10.2.1.3).

10.2.1.1 Full Tactic Information, No Future Offers

We start by discussing a setting where we have full tactic information (the controller knows with 100% accuracy what negotiation tactic each and every seller will use if negotiated with) but no information about future offer. In other words, the buyer will commit itself as soon as a contract that ensures positive expected utility for it can be found.[16]

We will first be interested in the relative performances of the advanced concurrency strategies (*Analytic*, *Adaptive*) against the best of the *Simple* strategies in each setting. The main advantage of the advanced concurrency strategies in this setting is that they can vary the number of negotiations to have according to the situation, whereas the *Simple* strategies will always have the same number of negotiations. Both advanced strategies consider not only the possible fees associated with extra contracts, but also the expected utilities of different negotiations. It might be a good idea to avoid negotiating with bad or mediocre opponents when there are also good ones around. The *Adaptive* concurrency strategy is of course much better at this than the *Analytic*, but also the analytic can avoid bad opponents in situations where it has a few good opponents and then a bad one. It simply negotiates only with the good ones. The *Simple* strategies are unable to attain such flexibility.

However, this flexibility can also be counterproductive. This can occur when the first negotiation providing a positive utility will be followed through, because it can mean that the buyer using the advanced concurrency strategies is able to find a contract earlier than the buyer using the *Simple* strategy with only a few (or even just one) negotiations. For example, if the fifth opponent in the negotiation queue would be able to provide positive utility, the *Simple* strategies with $1-4$ negotiations will not see this fifth opponent now (because they only negotiate with

[16] This is similar to the situation we had in the market setting in chapter 6.

up to four opponents), but might encounter it or some similar provider in the next matching or later. And in some situations this may be problematic because the expected utility usually increases over time (if the buyer can be affected), and therefore any delay that the *Simple* strategies may encounter, can work in their favour (they get better expected utility when they finally find a contract).

We discuss first the case where the *Adaptive Counter* tactic is used at the `Negotiator` level. We start from the cases where the *Random* opponent selection is used. Here, the differences in expected utilities among the sellers that the buyer encounters are at their largest because the opponents vary greatly in both quality and negotiation tactic.[17] With *Constant 0.0*, the advantage of the *Analytic* strategy over the *Simple* strategies is based on the fact that it can take into consideration the expected utilities of different negotiations and therefore can choose the number of negotiations so that the expeced utility is maximised. This may mean, for example, negotiating with only the first four opponents because the fifth opponent provides very low quality and we run a risk of getting into a contract with that opponent if we negotiate with it (especially if we expect it to use a negotiation tactic that will lead to a contract quickly, like the *Pure Behavioural* one). The *Analytic* strategy can, therefore, avoid bad contracts whereas the *Simple* strategies take whatever they get which in this setting can be almost anything, because the *Adaptive Counter* succeeds basically against any opponent (because all opponents provide a positive expected utility in this setting). This also means that all concurrency strategies (also the *Simple* ones) will get a contract very quickly, so there is no significant difference between the advanced and the *Simple* strategies in this respect. In addition, the *Adaptive* concurrency strategy can do even better than *Analytic* strategy, because it can avoid contracts with bad and mediocre opponents much more efficiently than the other strategies.

When the *Constant 1.0* decommitment policy is used, negative expected utilities are also possible. This is so especially in cases where only the buyer is affected and will mean that all strategies have to wait until a positive expected utility can be reached. Here, however, the advanced concurrency strategy may be able to find opponents providing it with a positive expected utility earlier than other strategies. Due to the *Random* opponent selection, the first such opponents may not appear in the top of the negotiation queue and therefore the *Simple* strategies with a small number of negotiations might not spot them when they first arrive. This

[17]Very low quality providers may provide utilities that are close to zero, whereas the high quality providers using the *Exponential Time-Dependent* tactic might be able to offer the buyer a utility of almost 0.50.

may mean that the *Simple* strategies actually do better than the advanced in some cases. This is because the later contracts usually provide a higher expected utility (the risk of disastrous effect is lower). On the other hand, the advanced strategies can avoid the mediocre contracts and extra decommitments just as in the *Constant 0.0* case. Because of the high decommitment cost here, the *Analytic* strategy will be very careful to avoid extra contracts and will often negotiate very little beyond the first applicable opponent, which means it will avoid decommitments but get into contracts earlier. The *Adaptive* strategy can avoid the bad and mediocre and needs to be less worried about the extra contracts, so it should perform well with low effect probabilities (when the effect probability does not make contracts happen that much later even for the *Simple* strategies).

Of course when only the seller is affected, getting into contracts faster is actually an advantage because it maximises the probability of very lucrative seller failure. Here, any contract is acceptable already in the first negotiations, so also the *Simple* strategies will get a contract quickly. However, because of the high decommitment fee, all the best strategies except the *Adaptive* are going to have very few negotiations, which means they will not encounter that many opponents and the results may not be great. The *Adaptive* strategy, on the other hand, can explore a wider range of opponents and is probably going to do better than the competition. The *Analytic* is probably less successful here. When both are affected, these two effects are combined, but because the risk of their own effect is neutralised by the chance of the seller effect (both being equally likely), the situation is more similar to the case where the seller is affected, because the expected utilities are going to be non-negative from the start. So, the *Adaptive* is expected to do well and the *Analytic* is expected mostly to be unable to beat the best *Simple* strategies.

Now, when the *Quality* opponent selection is used, the situation stays much the same. The *Adaptive Counter* will still have big differences in the expected utility among the opponents it encounters, because although everybody it encounters provides a high quality service, they can still use different negotiation tactics and the *Adaptive Counter* adapts to each of these tactics in a different way. There is quite big a difference whether or not it negotiates with somebody using the *Pure Behavioural* tactic or somebody with the *Exponential Time-Dependent* tactic, although when all sellers the buyer encounters are high quality, the differences are smaller than with the *Random* opponent selection.[18] This means that the advantage of the *Analytic* strategy diminishes also in the *Constant 0.0* cases because even the worst options are still decent and it will often be unable to do much better. However, the *Adaptive* strategy will still be going strong, because it can choose among the best.

The cases where the *Constant 1.0* decommitment policy is used will go much like with the *Random* opponent selection. However, because all providers are high quality, it means that the advanced and the *Simple* strategies find their contracts closer together and therefore the negative effects of the flexibility are probably going to be milder in the cases where the buyer is affected. Other cases are likely to remain much the same, although the *Adaptive* strategy's advantage is slightly decreased because all opponents produce good results when successful.

When the *Expected Utility* opponent selection is used, there is very little difference between the best strategies. This is because the best opponent is always going to be the first and the only opponent the sellers will negotiate with, so all strategies find their opponents around the same time and often with the same opponent.[19]

There is one more observation to make before our first hypothesis in this chapter and this is to remark that the way we calculate expected utilities in some parts of our model mean that the expected utilities are sometimes slightly too high. This is because they estimate the contract times of the seller slightly optimistically in the negotiations that follow the last matching. We have taken a somewhat simplified approach in a couple of places and this means that the estimates are not always exactly accurate. The difference to the actual values is not great, but it may be detectable in certain settings.

The two inaccuracies are very similar in nature and affect the way the expected utilities are calculated especially in the last matching. First, when we calculate expected utilities for a given negotiation we use the average of first and last possible contract time to estimate the expected utility. Here, we do not consider the contract time distributions we use in the concurrency strategies at all. This method does not take into account the fact that in the last negotiation the sellers will have their deadlines earlier and therefore the results, if any, will also occur earlier. This

[18] With an opponent of quality 0.99999, the *Pure Behavioural* tactic will yield a utility of around 0.25, whereas an opponent with the same utility but using the *Exponential Time-Dependent* will give almost 0.50.

[19] The *Simple* strategies may, in some cases, waste an opponent or two when the *Constant 1.0* policy is in force because they will have to negotiate also when they have no chance of success. However, the differences are probably too small to be detectable over the noise.

problem is especially clear when the buyer's effect probabilities are very high, because then the effect probability decreases sharply during the last negotiations, which means that the effect probability is going to be higher on average than in our calculation. As a result, the expected utility is going to be slightly too high in these cases. This means that in some cases we might start negotiations that lead to a slightly negative expected utility.

The second inaccuracy has to do with how advanced concurrent strategies estimate the probability of extra contracts. Both the *Analytic* and *Adaptive* concurrency strategy use the contract time distributions as discussed in section 10.1.2.2, but we have simplified the calculation in some cases and, as with the other inaccuracy, this involves not taking into account the fact that in the last matching the negotiations are going to take a shorter time on average[20] and, therefore, for example the random strategies will lower their target utility levels to ensure a reasonable chance of success.[21] This means that the probabilities of early contracts are sometimes too low and this may lead the advanced concurrency strategies to underestimate the risk of too many contracts and have too many negotiations.

The effect of these two inaccuracies are limited mostly to the cases where the buyer is affected with very high probability and the *Constant 1.0* decommitment policy is used (so that decommitment is costly). In other cases, their effect is going to be too small to be noticable. However, where it does apply, it will be around in all our settings and, in the later settings, we will just refer to this explanation. It is also worth noting that the different negotiation tactics will be affected to a different degree. The worst affected are the random tactics (*Random* and *Random Behavioural*) because both inaccuracies will influence the calculation of expected utilities and probabilities of multiple contracts. The *Exponential Time-Dependent* tactic will be among the least affected with the first inaccuracy because the negotiation will often end very late (often near the seller's deadline) and this means that the error will be smaller than in other cases. The second inaccuracy may affect things, especially if some of the sellers involved use the *Random* tactics because the negotiations will not often take until the last possible turn but may end earlier increasing the risk of simultaneous contract with the *Random* tactic.

Now, because these inaccuracies affect the *Random* negotiation tactic more than they do the *Exponential Time-Dependent* one, the *Adaptive Counter* tactic is

[20]This is because the sellers will almost always have a deadline before the last possible deadline which is just before the delivery time.

[21]In other words, the counter tactics do take this effect into account when calculating the optimal offers, it just is not considered when expected utilities are considered.

going to be less affected especially with sophisticated opponent selection because it allows the buyer to choose mostly sellers using the *Exponential Time-Dependent* tactic. The buyer using the *Random Counter* negotiation tactic will not be able to make this distinction and it will meet opponents using both tactics. This means it will be less effective.

We are then ready to contend:

Hypothesis 22. When there is full tactic information, but future offers are not considered, and the *Adaptive Counter* tactic is used at the `Negotiator` level:

a. the *Analytic* strategy will be able to beat all the *Simple* strategies in (almost) all cases, when the *Constant 0.0* decommitment policy and *Random* opponent selection are used and, in some cases, when the *Quality* opponent selection is used. It can lose to the best of the *Simple* strategies in some cases with the *Constant 1.0* decommitment policy when the buyer can be adversely affected, especially under the *Random* opponent selection but possibly also under *Quality* opponent selection.

b. the *Adaptive* strategy will be able to beat all the *Simple* strategies in most settings with the *Random* or *Quality* opponent selection. It may, however, lose to the best *Simple* strategies when only the buyer is affected and the *Constant 1.0* decommitment fee is used.

c. the *Adaptive* strategy will be able to beat the *Analytic* strategy in most settings with the *Random* or *Quality* opponent selection.

Now, the *Random Counter* tactic will fail all negotiations with the sellers using behavioural negotiation tactics (so, 50% of the time in our setting) and where successful, it will always get the same utility from a provider of the same quality because it always makes its offers based on the opponent quality (and expected negotiation time) alone. This makes the situation quite different from what we just discussed. The *Random Counter* tactic experiences less extra contracts than the *Adaptive Counter* one, which benefits both the best *Simple* strategies and the *Analytic* strategy because they can negotiate more and have a better chance of finding good quality providers. However, the *Adaptive* concurrency strategy loses some of its power, because with the *Random Counter*, it will be unable to exploit the sellers using the *Exponential Time-Dependent* tactic to quite the same level as before. Moreover, whereas the *Adaptive Counter* exploiting the *Exponential Time-Dependent* tactic can be certain that no opponent of a similar quality or lower will be able to match that expected utility (hence, no chance of extra contracts), this does not hold with *Random Counter*. The offer made is acceptable to both opponents using the *Random* and the *Exponential Time-Dependent* tactic with similar quality or even those with a slightly lower quality.[22] This means that the *Adaptive* strategy must be slightly more careful to avoid extra contracts where necessary. The difference is not huge, since also *Adaptive Counter* may sometimes have to rely on sellers using the *Random* tactic, but it is there, nevertheless.

These differences are not obvious when the *Random* opponent selection and the *Constant 0.0* decommitment policy are used. The advanced strategies have a clear advantage over the *Simple* strategies because of their flexibility and the *Adaptive* strategy has an advantage over the *Analytic* strategy, because it will be able to negotiate with more opponents and ignore the bad and mediocre ones (it still can do that). When the *Constant 1.0* policy is in force, the first differences begin to show. Because the *Simple* strategies with the low number of negotiations may take a while before they find a contract (whereas the advanced strategies are able to find the few opponents that they can succeed with), the advanced strategies can often find contracts quicker. This means more cases where the *Simple* strategies can beat the advanced strategies and, on the other hand, less chances for the *Adaptive* concurrency strategy to shine. They will, of course, have no trouble beating the *Simple* strategies when early results are preferred (the seller is affected) and here, both advanced strategies are likely to do well. Also here the cases where the buyer is affected with very high probability, the inaccurate expected utility and contract time estimates may cause a problem as discussed above and as also discussed above, the *Random Counter* is more susceptible to these problems and may therefore have more trouble in these cases.

The differences become very clear when we consider the case with the *Quality* opponent selection. Here, all opponents are high quality, so all successful negotiations started at around the same time will lead to a very similar expected utility. This, together with the fact that the *Quality* opponent selection makes it less worthwhile to have a large number of negotiations (chances of finding better contracts are not large enough to counter the increased risk of extra contracts), and intensifies the negative effect of the advanced strategies' flexibility. When the *Constant 0.0* decommitment policy is used, this means that the advanced strategies are almost always able to find an opponent that gives them positive utility,

[22]This is because the *Random Counter* tactic does leave some utility to the opponent too.

whereas the *Simple* strategies with a low number of negotiations have a reasonable chance of not finding one in the first matching (because the opponent they encounter might use a behavioural tactic).[23] This means that in cases where the contract time matters (the seller or both are affected), the best *Simple* strategies are probably going to outperform the advanced ones. They enter into contracts later and have a lower chance of adverse effect removing the contract. The *Simple* strategies may also have a slight advantage also when only the buyer is affected, because the first opponent using the non-behavioural negotiation tactic may have a slightly higher quality in the later matchings, because of the later entries.[24]

As in the case with the *Constant 1.0* decommitment fee, the advanced strategies are still going to find contracts slightly quicker than the best *Simple* strategies. This means that where the buyer is affected, the advanced strategies are more likely to be in trouble. On the other hand, when the seller is affected, getting into contracts quickly is still a good thing, especially when the probability of the seller effect is very high, so we would expect that the advanced strategies do well when the seller is affected (especially when the effect probability is high so that early results are strongly preferred).

With both *Quality* and *Expected Utility* opponent selection, the *Adaptive* concurrency strategy may have an advantage over the *Analytic* concurrency strategies in cases with very low effect probabilities and the *Constant 0.0* decommitment policy. This is because with the *Random Counter* negotiation tactic, the chances of success are very high, but less than one. This means that there may be cases where the buyer might be able to get a slightly higher expected utility out of two or more opponents than with only one. This is because the probability of success will decrease with an increase of expected utility more slowly with multiple opponents than with one opponent. An example might be useful here. So, if we have a probability of success of say 99.5% with one seller, but if we have two very similar opponents with that success probability, we would have a probability 99.9975% that at least one of the negotiations will be successful. This means that we can increase the expected utility we demand and still have a very reasonable probability of success. However, this is not always possible. For starters, we run a risk of having two contracts and with a positive decommitment fee, this strategy may

[23]So, for the *Simple 1* strategy, finding an opponent that uses a non-behavioural tactic at first matching, the probability is roughly 0.5. For an advanced strategy, the probability that none of the ten opponents it considers uses a non-behavioural tactic is minuscule, $0.5^{10} = 0.0977\%$.

[24]There is, on expectation, only 25 sellers using non-behavioural tactics in the market in the first matching and around 100 more are expected to arrive at some point later. The chances are that one of these later arrivals has a higher quality than the best early bird.

fast become unusable. On the other hand, we need sellers that have very similar qualities and are using the *Random* negotiation tactic [25] to pull this off and given that most contracts are formed very quickly after the start of the experiment, there is a limited space for success. On the other hand, if we have a positive chance of adverse effect, the expected utility increase of a small price adjustment might not increase our expected utility in a significant way even when all the other requirements are fulfilled.[26] Therefore, any difference will likely to be limited in cases with very low effect probabilities.

Given all this, we contend:

> **Hypothesis 23.** When there is full tactic information, but future offers are not considered, and the *Random Counter* tactic is used at the `Negotiator` level:
>
> a. the *Analytic* strategy will be able to beat all the *Simple* strategies in (almost) all cases, when the *Constant 0.0* decommitment policy and *Random* opponent selection are used and in some cases (when only the seller is affected), when the *Constant 1.0* policy is used. It can lose to the best of the *Simple* strategies in some cases with the *Constant 1.0* decommitment policy when the buyer or both parties can be affected, with both the *Random* and *Quality* opponent selection and with the *Constant 0.0* decommitment policy (especially when the seller or both parties can be affected) when the *Quality* opponent selection is used.
>
> b. the *Adaptive* strategy will be able to beat all the *Simple* strategies in most settings with the *Random* opponent selection. It may, however, lose to the best *Simple* strategies when only the buyer or when both parties can be affected when the *Constant 1.0* decommitment fee is used with both the *Random* and the *Quality* opponent selection. It will also lose to the best *Simple* strategies when a *Constant 0.0* decommitment policy and the *Quality* opponent selection is used and the seller or both parties can be affected.

[25] The *Expected Time-Dependent* tactic is not susceptible to this trick, because any lower price will never be acceptable to the seller and therefore increasing the expected utility by lowering the price will not work. With the *Random* tactic, there is a small difference between the optimal price in one negotiation and the reservation price, so a slight increase may be possible. However, the probability of success decreases very rapidly with the price adjustment and the room for improvement is limited.

[26] This problem is especially clear here, since with no future offers, the contracts are formed very quickly which means that the probabilities of effect are at their highest.

c. the *Adaptive* strategy will be able to beat the *Analytic* strategy in most settings with the *Random* opponent selection. With other opponent selection strategies, the *Adaptive* strategy's advantage is limited to the cases with the *Constant 0.0* decommitment policy and very low probability of adverse effect. There may also be cases with the *Constant 1.0* decommitment fee, where the *Adaptive* gets beaten by the *Analytic* strategy.

10.2.1.2 Full Tactic Information with Future Offers

When we take into consideration the future offers, the situation changes. Here, it becomes very important what sort of opponents (the number plus the quality and tactic distribution) the concurrency strategies will encounter during the run. In particular, we are not after the first contract that yields positive expected utility, but the very best contract we can hope to find. As explained above, this means we use a utility threshold which typically decreases as the time goes on so a contract that would not have been acceptable earlier might be so later. This is because the probability of finding a better one has decreased. Now, the different concurrency strategies obviously will encounter a different selection of opponents during a run. For example, *Simple 1* will only see 10 opponents at most, whereas *Simple 10* might see as many as 100. Moreover, the opponent selection strategies will have a clear impact on how good those opponents encountered are.

There is also a clear difference between the advanced and *Simple* strategies. The *Simple* strategies never encounter the same opponent twice. If they see it, they will negotiate with it and then they can never negotiate with that same opponent again. In contrast, the advanced strategies might encounter some of the opponents over and over again, because they often do not negotiate with the opponents they encounter. This can be both an advantage and a problem. It can be an advantage because if an advanced strategy sees a good opponent too early for that opponent to be useful, it can wait for a few matchings and negotiate with it when the probability of adverse effect has decreased to a more acceptable level (assuming the opponent in question is still in the market, of course). However, it can also be a problem. When the *Random* opponent selection is used, the best approach for finding the best opponents may in some cases be meeting as many opponents as possible. The difference between the *Simple* and the advanced concurrency strategies is not as great as it may seem at first because the opponents the buyer sees are selected at *Random* at every matching, so it is likely that most of the opponents are going to be 'new' also for the advanced strategies. Also advanced strategies will be able compensate with their ability to choose the best number of opponents to actually negotiate with. However, there is room for the simple strategies to do very well against the advanced strategies in cases with low probabilities of buyer effect and the *Random* opponent selection. With medium to high effect probabilities, with *Constant 1.0* and more advanced opponent selection methods, the positive effect of waiting to negotiate will be dominant and it will not matter that the advanced concurrency strategies see less of the opponent pool, because the *Simple* strategies will be unable to use the opponents they see too early.

Now, an interesting observation which holds generally is that when only the buyer is affected, the buyer cannot influence the probability of its adverse effect and this probability is independent on with which opponent and when a possible contract is formed. This means that if the probability of adverse effect is 0.8, then 80% of the runs end up having an adverse effect at some point no matter what the buyer does. The only thing the buyer can influence with his behaviour is whether or not he is in a contract if and when the adverse effect occurs. If he is careful, he might not be and if he is not that careful, he might be. Now, in some settings, both of these cases produce the same outcome. For example, when the fee is zero, having no contract when the adverse effect occurs produces zero utility, but if it is free to get rid of any contract so does the situation where the buyer has a contract. Therefore there is no need to be careful. Of course should the fee be positive, the situation is very different. Then having no contract still produces the utility of zero, but having a contract to decommit means negative utility. This means that being careful is useful and when the fee is very high, it is crucial to avoid getting into contracts when an adverse effect is likely. What all this means is that when the fee is zero, the only thing that matters is how good are the contracts that will actually be performed and here the advanced strategies may have only a limited edge over simple strategies and that edge can only be based on that fact. This may not always be enough.

In contrast to this, when it comes to the cases where the seller is affected, waiting for a while will actually decrease the probability of failure, because some of the opponents that would have failed exit the market and therefore the probability that the seller we enter into a contract with is going to be affected is lower.[27] On the other hand, sometimes when a decommitment is actually better than a

[27]One way to put this difference is that the buyer can choose its contract partner, but it is unable to change itself.

performance (over-compensatory decommitment policy), it is useful to enter into a contract as soon as possible to maximise the probability of a seller failure. When the fee is one, the best approach is to get into a contract as soon as possible and the difference between the advanced and the simple concurrency strategies depends solely on how well they will be able to enter into a contract during the first negotiations and how good the resulting contract is. The advanced strategies are likely to have an edge when opponents are chosen at random, because they can probably find better contracts than simple strategies with low numbers of negotiations. On the other hand, when opponent selection is more sophisticated, also the simple strategies will be able to find good opponents to negotiate with and the advantage of the advanced strategies wanes.

Also, as we discussed in section 10.2.1.1, our expected utility estimates may be slightly too high in the last matching, especially if multiple negotiations are considered, the *Constant 1.0* decommitment policy is used and the probability of buyer effect is very high. The situation here is not any different from the case with the future offers, because all the contracts will be formed very close to the deadline (after the last matching) and the future offers are not considered there.

Otherwise, the difference between the *Adaptive Counter* and the *Random Counter* tactics is in this setting somewhat smaller than it was without the future offers. That is because the future offers remove much of the bad and mediocre opponents from consideration in the early rounds in both cases. Therefore the difference between the two has more to do with the *Adaptive Counter*'s ability to exploit the *Exponential Time-Dependent* tactic better. Of course if and when the contracts are formed at the very last minute, the differences discussed in the previous subsection have a role to play because the future offers are no longer considered in the last matching. Moreover, sometimes lower threshold levels produce better results, which means that there may be more than one opponent available also in the earlier matchings. Here, there is a clear difference between the two negotiation tactics, because the probability of extra contracts (and contracts in general because some negotiations will fail) is smaller with the *Random Counter* tactic than it is with the *Adaptive Counter* tactic. This means that the *Simple* strategies with a higher number of negotiations can often be used. This could be useful for example when the *Constant 1.0* decommitment policy is used and both parties can be affected. Here, the only problem with the high fee is the extra contracts and that worry is relatively small with the *Random Counter* tactic, which means that the *Simple* strategies may be able to do well. With the *Adaptive Counter* tactic the risk of extra contracts is more significant and the *Simple* strategies may have to have

have a smaller number of negotiations to avoid extra contracts and the advanced strategies may be able to outperform the *Simple* strategies in most cases. Also the advanced strategies will be able to find sellers that use the *Exponential Time-Dependent* tactics more often and exploit them, whereas the *Random Counter* is unable to do that. However, the main advantage of the advanced strategies is going to be the ability to wait for the good opportunities without negotiating, however, and this is likely to be much more important in practice.

Given all this, we contend:

> **Hypothesis 24.** When there is full tactic information and future offers are considered, the advanced concurrency strategies (*Analytic*, *Adaptive*) perform at least as well as the best *Simple* strategies in most settings, but there may also be settings where the best *Simple* strategies outperform the advanced strategies. We expect the following patterns to emerge:
>
> a. When only the buyer is affected and the *Constant 0.0* policy is used, the difference between the best *Simple* strategies and the advanced concurrency strategies are likely to be small or even missing.
>
> b. When only the seller is affected and the *Constant 1.0* policy is in force, the advantage of the advanced concurrency strategies is likely to be the weaker the more sophisticated the opponent selection is. We expect (almost) no difference when the *Expected Utility* opponent selection is used.
>
> c. When both parties can be affected and the *Constant 1.0* policy is used, there may be effectively no difference when the *Random Counter* negotiation tactic is used at the `Negotiator` level.
>
> d. There may be cases with the *Random* opponent selection and the *Constant 0.0* decommitment fee and low effect probability where the *Simple* strategies outperform the advanced strategies. There may also be cases with the *Constant 1.0* decommitment policy when only the buyer can be affected with very high probability (like 0.9) where the advanced strategies are beaten by the *Simple 1* policy.

The next question for us to ponder is the relative performance of the advanced concurrency strategies. Although the future offers remove many mediocre and bad opponents, they do not do it with a perfect accuracy, especially not near the deadline, since the threshold drops whereas the expected utility increases when time progresses. Because we use averages in our calculations, it means that there are always cases where the thresholds are not met in the earlier matchings and the contracts are formed nearer the deadline, whereas as we just mentioned, the thresholds are lower or even zero in the last matching. This means that the *Adaptive* strategy will be able to do relatively well. The effect is likely to be at its strongest when the differences in utility are large (*Random* opponent selection for both and *Quality* opponent selection for the *Adaptive Counter* tactic) and where waiting improves the outcome for the buyer and with the future offers it may well be that the buyer will often have to wait until the last matching to reach that best deal. This means cases with only the seller affected (Constant 0.0) and the buyer affected (Constant 1.0). The effect can also be there (albeit weaker) when both parties are affected (Constant 0.0). Thus, we say:

> **Hypothesis 25.** With full tactic information and future offers, the *Adaptive* concurrency strategy is always at least as good as the *Analytic* concurrency strategy and it beats the *Analytic* strategy when only the seller is affected (Constant 0.0) and when the buyer is affected (Constant 1.0) with the *Random* opponent selection. With the *Adaptive Counter* tactic used in the `Negotiator` level, the *Adaptive* strategy is also able to beat the *Analytic* one in the same cases with the *Quality* opponent selection and with cases where both are affected (Constant 0.0) with the *Random* opponent selection.

The third topic we discuss is the role of the opponent selection strategies. We believe that the move from the *Random* opponent selection to the *Quality* opponent selection is likely to improve performance in most settings and in most cases this is likely to be useful. This is because the change improves the type of opponents the strategies encounter. Specifically, the *Quality* opponent selection means meeting mostly high quality providers and since the high quality providers usually offer the buyer a higher expected utility than low or mediocre quality providers, the change usually improves the buyer's expected utility. On the other hand, the *Quality* opponent selection also means that finding good opponents requires less negotiations and therefore strategies with only a few negotiations do better and the risk of multiple contracts decreases.

When it comes to the switch from the *Quality* to the *Expected Utility* opponent selection, however, the picture is less clear. With the *Random Counter* negotiation tactic, it seems clear that the effect to the advanced strategies should be minimal. The *Quality* opponent selection already orders the opponents that can provide us with a success in descending order of expected utility and the only thing the *Expected Utility* opponent selection does is to remove the opponents that the buyer cannot succeed with. However, the advanced strategies are usually able to find these opponents in any case, since it is likely that at least one of them is going to be among the ten opponents it can negotiate with. The *Simple* strategies, in contrast, may well find that the *Expected Utility* opponent selection works for them, because it will ensure that the first opponent in the list is always the best one in the market, whereas with the *Quality* opponent selection it could also be an opponent that the buyer could not succeed with. So, there should be some improvement in some cases there.

With the *Adaptive Counter* tactic at the `Negotiator` level, also the advanced strategies, especially the *Analytic* concurrency strategy, may find that the *Expected Utility* opponent selection improves their performance. This is because it will remove the opponents using the behavioural strategies mostly from contention also in the last matching and, therefore, improve the selection for the advanced strategies. The *Adaptive* strategy is likely to be less affected because it is better able to pick these opportunities also from a bit further back in the list, whereas the *Analytic* strategy may find this difficult. This effect is especially likely with the buyer being affected and the fee being high (Constant 1.0). For the *Simple* strategies, the move from *Quality* to *Expected Utility* may not be all that good in all cases. In most cases it means that it will enter into contracts quite late in the game and that means that it has negotiated with many of the best opponents that market had to offer, so although it can probably find a decent contract in the end, it may not be the best the market had to offer. In contrast, with the *Quality* opponent selection, some of the opponents encountered are not so good, but it is also tricky to find that best deal at the last minute, because that best quality provider in the last matching might be using a behavioural tactic. So, it seems likely that this change is most advantageous for the *Simple* strategies when the expected utility increases rapidly in the last minute, so when only the buyer is affected and the fee is one (*Constant 1.0* policy in use). Thus, we contend:

Hypothesis 26. With the full tactic information and future offers, the change from the *Random* to *Quality* opponent selection is likely to be useful in most settings with both the *Adaptive Counter* and the *Random Counter* tactics. With the *Random Counter* tactic, only

> the *Simple* strategies will benefit from moving from the *Quality* to the *Expected Utility* opponent selection. With the *Adaptive Counter* tactic, the best of the *Simple* strategies and the *Analytic* strategy are likely to find this advantage in cases where last minute deals offer good utility (the seller is affected (*Constant 0.0*) and especially when the buyer is affected and the *Constant 1.0* decommitment policy is in force).

Of course the whole point of using the future offers was that we expected them to improve the expected utility for the buyer. Since one of the possible values for the expected future utility is zero, we can be assured that taking future offers into account will always be at least as good an option as not taking them into account. However, the question is then whether or not taking the future offers into account *improves* the expected utility. This seems likely in many circumstances because the future offers remove bad and mediocre opponents from consideration and allow the strategies to concentrate on finding the best opponents they can.

The future offers work a bit like the *Adaptive* concurrency strategy in this regard. The opponents failing to provide a certain level of expected utility are ignored (any negotiations with them are made to fail on purpose). It will of course help also the *Adaptive* concurrency strategy because it will make it consider not only the opponents available now but also the situation later. This means that the *Adaptive* concurrency strategy does not take the first contract it sees, but can wait for a better one.

However, there may well be cases, where considering future offers does not help. One such case is the situation where the positive expected utility is difficult to come by. So, when the buyer effect probability is very high, getting any contract to give a positive expected utility might be so rare an occurence that any such contract should just be taken. On the other hand, in situations where getting into a contract early is the best approach, like when only the seller is affected and the fee is 1.0, the best contracts are usually made in the first matching, which means that any threshold is going to be quite low and only remove a handful of opponents at most. The effect is likely to be quite small. The third case concerns situations where only the buyer is affected and the fee is zero (the waiting does not improve the probability of success and failure to get a contract and decommitment produce the same result, zero utility) and advanced (not *Random*) opponent selection is used (the differences between different options are relatively small). This means that there is little room for improvement after finding the first deal. Specifically, we say:

Hypothesis 27. When tactic information is available, taking future offers into account improves the utility in almost all cases. However, there are exceptions that include:

a. when getting a positive utility at all is not certain (very high effect probability for the buyer or both parties, especially with the *Constant 1.0* decommitment policy),
b. when getting into a deal quickly is the best approach (only the seller is affected and the *Constant 1.0* policy is used) and a good opponent selection policy is used (the *Quality* or *Expected Utility* for the *Random Counter* tactic and the *Expected Utility* for the *Adaptive Counter*),

And finally, we discuss the differences between the two negotiation tactics: *Adaptive Counter* and *Random Counter*. The main differences between the two are, as has been discussed already, the fact that the *Random Counter* is unable to succeed with opponents using either of the behavioural tactics and that it is unable to distinguish between the *Exponential Time-Dependent* and the *Random* tactic but it makes the same offers against both these tactics. The *Adaptive Counter* can get a result with any opponent and when opponent tactics are known, it is able to distinguish between the two non-behavioural strategies. Now, the future offers remove many of the opponents using behavioural tactics from the scene, so the first difference between the two tactics therefore loses some of its effect and, on the other hand, the second difference still remains and actually may even get stronger, because the *Adaptive Counter* can try to target the *Exponential Time-Dependent* tactic especially when the more advanced (especially the *Expected Utility*) opponent selection methods are used.

However, when the *Random* opponent selection is used and the *Constant 1.0* fee is used, there is room for the *Random Counter* tactic to shine. This is mostly because of its lower tendency to get into extra contracts even when negotiating concurrently. It means that it can negotiate with more opponents and find better deals. However, because the advanced concurrency strategies are able to manage the risk of extra contracts, they are likely to be unaffected by this and, therefore, with the advanced strategies, the *Adaptive Counter* is likely to remain the superior tactic also in these circumstances. In particular, we contend:

Hypothesis 28. When tactic information is always available, the buyers using the *Adaptive Counter* tactic often outperform those that use the *Random Counter* tactic at the `Negotiator` level. However, there are also cases where *Random Counter* tactic does better and cases where their performance is the same. We expect the following:

a. the *Adaptive Counter* strategy's gets stronger when the sophistication of the opponent selection improves (we expect it to beat the *Random Counter* more often in cases using the *Quality* and *Expected Utility* opponent selection. With the *Adaptive* concurrency strategy, the *Adaptive Counter* is often a better choice than the *Random Counter*.

b. the *Random Counter* tactic is able to do better with the best *Simple* concurrency strategies when the *Constant 1.0* decommitment policy is used.

10.2.1.3 No Tactic Information with Future Offers

In this setting we remove the tactic information, so the buyer agent does not know what negotiation tactic each individual seller uses any more, but it only knows the probabilities of each tactic in the whole seller population. The advanced concurrency strategies use the information that can be gained from the opponent's tactic and, as discussed in section 10.1.2.2, the *Adaptive* concurrency strategy in the form we implemented it is not readily applicable to cases with no tactic information. Nevertheless, we can still use the *Analytic* concurrency strategy. Given that the *Analytic* uses the tactic information to estimate probabilities of simultaneous contracts and with the *Adaptive Counter* also in deciding what offer to make, also the *Analytic* concurrency tactic is likely to do less well without knowing the opponent's tactic.

One very interesting aspect is that there is likely to be a huge difference in how much the performance deteriorates with the loss of tactic information between the *Adaptive Counter* tactic and the *Random Counter* tactic. This difference was very clear in the *Negotiatior* level (see section 9.2.3) and adding more negotiations (especially when we cannot use the *Adaptive* concurrency strategy) will not help. In particular, the *Analytic* concurrency strategy is still bound to the suggestions from the `Negotiator` level and we know them to be often acceptable to all opponents (because that is how the *Adaptive Counter* tactic will react when it gets no tactic information) and therefore we would expect the performance to basically collapse with the *Adaptive Counter* tactic. On the other hand, the *Random Counter* tactic does this deterioration quite gracefully with slight or even no change in performance. Here, we will discuss only the case where there is no information and we discuss *Random Counter* and *Adaptive Counter* tactics separately.

We start from the case where the *Adaptive Counter* tactic is used at the *Negotiator* level. The whole tactic depends on choosing an optimal counter tactic to each and every negotiation. If we do not know what tactic the opponent is using, but only a probability distribution, this is obviously much more difficult. We can, and will, use the distribution of tactics, but as before, the *Adaptive Counter* will want to succeed in every negotiation because the risk of not succeeding is greater than the utility improvement offered by such risky offers. Of course given that there would be more than one contract available, the buyer should be more risk-seeking. However, the *Analytic* concurrency strategy cannot do that, because it is bound by the suggestions that individual `Negotiator`s give and these will not consider anything more than that one negotiation. We therefore contend:

> **Hypothesis 29.** When tactic information is not available, future offers are considered and the *Adaptive Counter* tactic is used at the `Negotiator` level,
>
> a. compared to the performance with full tactic information, the performance of the *Analytic* concurrency strategy with no tactic information is bad and the former beats the latter clearly in almost all cases. The possible exceptions include high probability of buyer effect when the *Constant 1.0* policy is used. There may also be a couple of cases, where the *Analytic* strategy performs better with the threshold levels gained under no tactic information. These are the cases with a high probability of buyer effect.
>
> b. the *Analytic* strategy will beat all the *Simple* strategies in some cases and it can, in turn, be beaten by the best *Simple* strategy when the probability of adverse effect is low and the *Random* opponent selection and the *Constant 0.0* policy are used.

The *Adaptive Counter* tactic would therefore need the *Adaptive* strategy, which could take into account the fact that there are more than one negotiation available. However, we can get some of this advantage by using the *Random Counter* tactic at the `Negotiator` level. Here, the advantage of *Random* counter tactic is that in the case of *Simple* concurrency strategies basically nothing changes: when we use the *Random* counter tactic, we assume that the opponent is using the *Random* negotiation tactic and do our adjusting accordingly by making an optimal offer given this belief. The fact that we do not know what the tactic really is (whereas before we did) makes no difference with the *Simple* concurrency strategies: We will still negotiate with that same fixed number of opponents, using exactly the same offers we would in the full information case and therefore, we get exactly the same results. However, for the *Analytic* concurrency strategy, something does change. Since we no longer know what tactics opponents use and have to rely on probability estimates, we are likely to make sub-optimal choices when it comes to choosing the right number of negotiations. Earlier we could see that the first three opponents are using tactics that we will not be successful with and could simply take the first four opponents to negotiate knowing we only have a chance to get one contract. Here, any of these four opponents could use a suitable strategy and we have to be more careful, especially if decommitting is expensive (the policy is *Constant 1.0*). This does not mean the *Analytic* concurrency control is going to be useless, but it does mean that it is going to be less effective and that there may be cases where it is actually worse than the best *Simple* strategies.

Hypothesis 30. When tactic information is not available, future offers are considered and the *Random Counter* tactic is used at the `Negotiator` level,

a. losing the tactic information has a relatively small effect on the performance of the *Analytic* concurrency strategy. However, the full tactic information is useful when the seller alone is affected, especially with the *Quality* opponent selection, or with low to medium probabilities of the buyer effect (with the *Constant 1.0* policy). There may also cases where the *Analytic* strategy performs better with the threshold levels gained under no tactic information.

b. the *Analytic* strategy will beat all the *Simple* strategies when the *Quality* opponent selection is used and waiting is useful. With the *Random* opponent selection, there may also be cases where the *Analytic* strategy beats the competition.

10.2.2 Experimental Setup

As in the earlier cases, we compare the performance of different strategies in cases where the market is entered 100 times (instead of just once), so we run the market separately with the different strategies for a hundred times, calculate the results of each strategy together, and repeat this processs 100 times to get an average performance and variance for hundred repetitions. So if we say one strategy beats another, it means that if strategies were to be used in a given market for 100 times, it is likely that there would be a difference between the two in the direction we described.

The results of this chapter will be somewhat different from any of the others in this thesis and that is mostly because we have quite a few different strategies, threshold levels and hypotheses. Specifically, we will only show a small number of these different things in our figures and we discuss only a small number of these cases. For example, in everything that follows in this chapter we always only show the results for the best of simple strategies, which means exactly what it sounds like, the best result any *Simple* strategy was able to obtain. We already discussed (section 10.1.3) the fact that we will have many different future offer thresholds but we only discuss the best of them. This means that we will not usually discuss which of the ten *Simple* strategies was the best in each and every case or what threshold levels worked best. This is because there are many other things to discuss too.

Instead of graphs and going through them in great detail, we use some figures, but mostly tables to summarise the results of our experiments and statistical tests. We will usually employ the same two-sided t-tests as before and, as we discussed earlier (section 10.1.3), we can have many different levels of future offer thresholds and we only represent one in our figures. However, whenever we say that a certain strategy beats another, this means that the best threshold level of that strategy beats all the threshold levels of the other in a statistically significant way and the p-value we discuss is the highest we found in any of these comparisons.

Because we will have many different cases, we will have also cases where a difference is found, although there might not be one,[28] and as we described in section 10.1.3, especially with future offers we might have different threshold levels none of which

[28] This is simple statistics. Statistics is never 100% accurate, but operates with probabilities. Even if we say there is only 0.1% chance that a certain difference is due to randomness, there is still that chance and given a thousand experiments, it would be expected to happen once on average.

is necessarily the optimal for a given strategy, and therefore differences might show in places where there really are no differences or fail to show in places where there are actual differences. The only approach to handle this is to consider that if we find or do not find a difference in one case, we would find or not find it in similar cases too. So, for example, we might not think too much of a single case of difference when only the buyer is affected with probability of 0.4, but if this happens also in cases $0.0-0.3$ and $0.5-0.9$, then we might consider the possibility that may be this difference is real and not just a coincidence. The same goes in the other direction. So if we have a clear difference in cases $0.0-0.3$ and $0.5-0.9$, but no difference at 0.4, we might consider this just a coincidence. Of course there might be something very peculiar happening at that particular setting, which would explain the difference (or lack thereof) but if no such explanation is forthcoming and also the problem case itself is only barely significant or unsignificant, this is unlikely.

10.2.3 Results

We will discuss the results in the same order we discussed the hypotheses above (section 10.2.1), so:

- Full tactic information, no future offers (section 10.2.3.1),
- Full tactic information with future offers (section 10.2.3.2) and
- No tactic information (section 10.2.3.3).

10.2.3.1 Full Tactic Information, No Future Offers

First, we consider the case where the buyer agent knows the negotiation tactic but will accept any contract that provides a positive expected utility. As with the hypotheses, we first consider the case where the *Adaptive Counter* tactic is used at the `Negotiator` level (hypothesis 22), then the case with the *Random Counter* tactic (hypothesis 23).

In figures 10.4–10.6, we have the performance of the different concurrency strategies with the *Adaptive Counter* negotiation tactic. From the figures, it seems clear that the *Adaptive* concurrency strategy does very well in most cases, but that the other strategies catch up when the opponent selection improves. The more detailed results are described in the table 10.1. The plusses and sometimes minuses describe the results of statistical testing, plusses meaning that the advanced concurrency strategy (on the top row, the *Analytic* and on the bottom row, the *Adaptive* strategy) has yielded results that are better (in a statistically significant way) than the best *Simple* strategies, and minus referring to the contrary situation (the *Simple* strategies fare better). The number of plusses or minuses gives the significance level, one plus means $p < 0.05$ level, two plusses $p < 0.01$ level and three the $p < 0.001$ level. A hyphen means that no statistically significant difference could be found. So, all this means that we can take a setting, let us say with the *Random* opponent selection and the buyer can be affected with a probability 0.7 and the decommitment policy is *Constant 1.0*. We can see that there is no difference (-) when it comes to the *Analytic* concurrency strategy and the best *Simple* strategies, but that there is a difference between the *Adaptive* and all *Simple* strategies in the former's advantage (three plusses, so at the $p < 0.001$ level). And to take another example, *Quality* opponent selection, buyer is affected with a probability of 0.5 and with the *Constant 1.0* decommitment policy, the best *Simple* strategy outperforms both the *Analytic* (two minuses so at the $p < 0.01$ level) and the *Adaptive* (one minus so at the $p < 0.05$ level).

From table 10.1, it is clear that the *Analytic* strategy beats the *Simple* strategies with the *Random* opponent selection and the *Constant 0.0* decommitment policy. There is also some difference with the *Quality* opponent selection with very low effect probabilities and when the seller is affected, also in other cases. The *Analytic* strategy is able to choose the number of negotiations flexibly which makes it possible to avoid some bad and mediocre contracts with the *Constant 0.0* decommitment policy. With the *Constant 1.0* policy, the *Analytic* strategy carefully avoids the very expensive extra contracts and is unable to beat the *Simple* strategies. However, with the buyer affected, it is sometimes able to find contracts earlier than the *Simple* strategies (it only needs to take the first opponent so the high fee is not a problem). This applies to cases where the buyer effect probabilities are 0.3, 0.4 and 0.8 with the *Random* opponent selection and 0.5 with the *Quality* opponent selection strategy. However, the cases with very high buyer effect probabilities (0.9 with the *Random* or the *Quality* opponent selection strategy) are associated with the expected utility estimate inaccuracies we discussed earlier. This means that the best *Simple 1* is often unable to find a good enough opponent to negotiate with in the last matching and will therefore rarely enter into counterproductive contracts (or any contracts at all), whereas the advanced strategies are more often able to find these problematic opponents that seem to

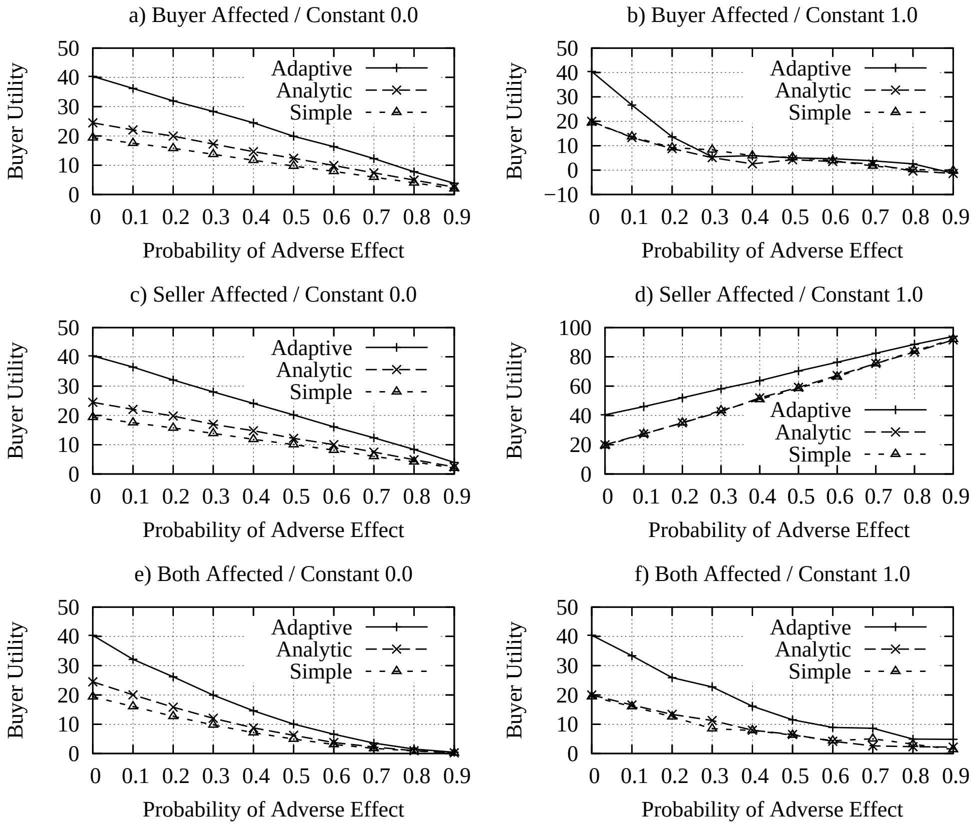

FIGURE 10.4: The performance of different concurrency strategies with the *Adaptive Counter* tactic and the *Random* opponent selection strategy (no future offers) (Hypothesis 22).

offer positive expected utility although in reality the expected utility is slightly negative. However, with the *Expected Utility* opponent selection, the buyer is able to find sellers using the *Exponential Time-Dependent* tactics and avoid most of the problems caused by inaccuracies: all the best strategies only negotiate with the first seller available in the final matching (no additional risk of extra contracts) and usually when a potential opponent is available, it is using the *Exponential Time-Dependent* tactic (decreasing the overestimation of expected utility). These findings are consistent with hypothesis 22.a and, therefore, we can accept that part.

Now, the *Adaptive* concurrency strategy is able to beat all the *Simple* strategies in all settings with the *Random* opponent selection and the *Constant 0.0* policy. It is also able to repeat this deed with one exception (both affected at probability of 0.9) when the *Quality* opponent selection is used. With the *Constant 1.0* policy, it has a bit more trouble, but it still beats the *Simple* strategies in most cases where the seller and both parties are affected (where the earlier contracts are an asset not a

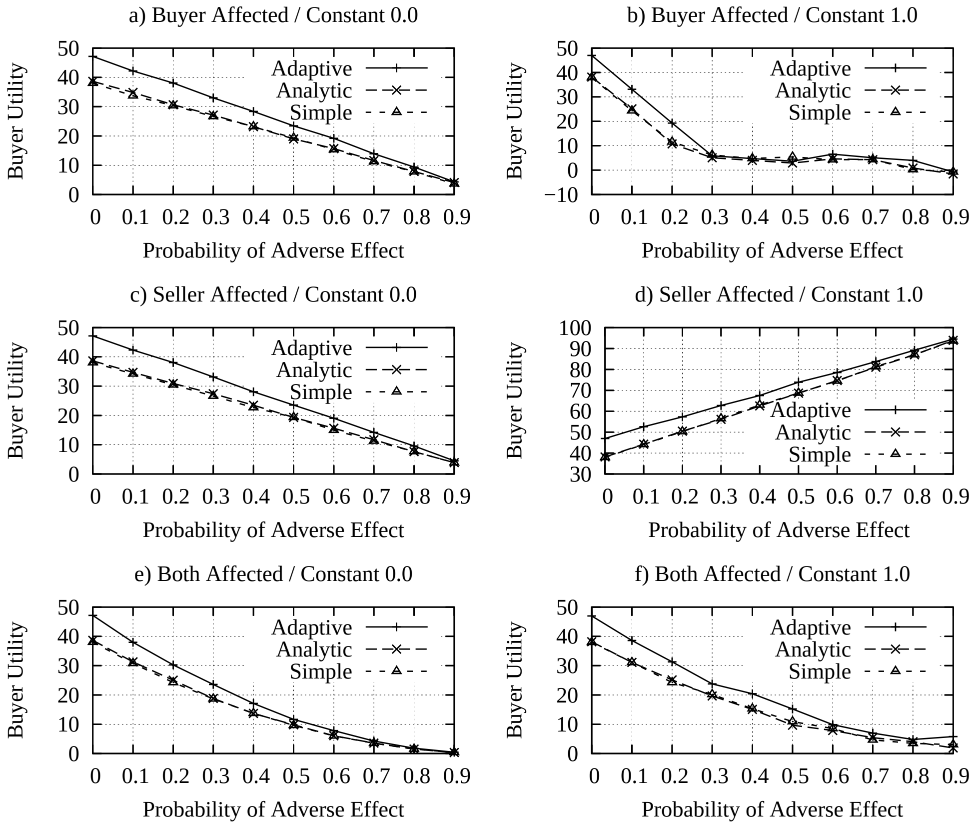

FIGURE 10.5: The performance of different concurrency strategies with the *Adaptive Counter* tactic and the *Quality* opponent selection strategy (no future offers) (Hypothesis 22).

problem) with both the *Random* and *Quality* opponent selection. When the buyer is affected, there are some cases where it still able to beat the *Simple* strategies but it will also lose in a couple of places (with the probability of 0.3 and 0.9 with *Random* and with the probability of 0.5 when the *Quality* opponent selection is used). The *Adaptive* strategy is able to avoid the bad and mediocre much better than *Analytic* and is therefore able to able to use the ability to negotiate with many opponents with limited chance of extra contracts to its advantage also with the *Constant 1.0* decommitment fee. This is also possible with the *Quality* opponent selection because that does not remove the variations in the expected utilities among the sellers it encounters. As for the *Analytic* strategy, it faces the problem of efficiency when the buyer is affected and the *Constant 1.0* fee is used. In some cases, it is able to find a contract earlier than the *Simple* strategies. The case with the buyer effect probability of 0.9 is analogous to the case with the the *Analytic* strategy. Everything described is consistent with hypothesis 22.b and, therefore, we can accept it.

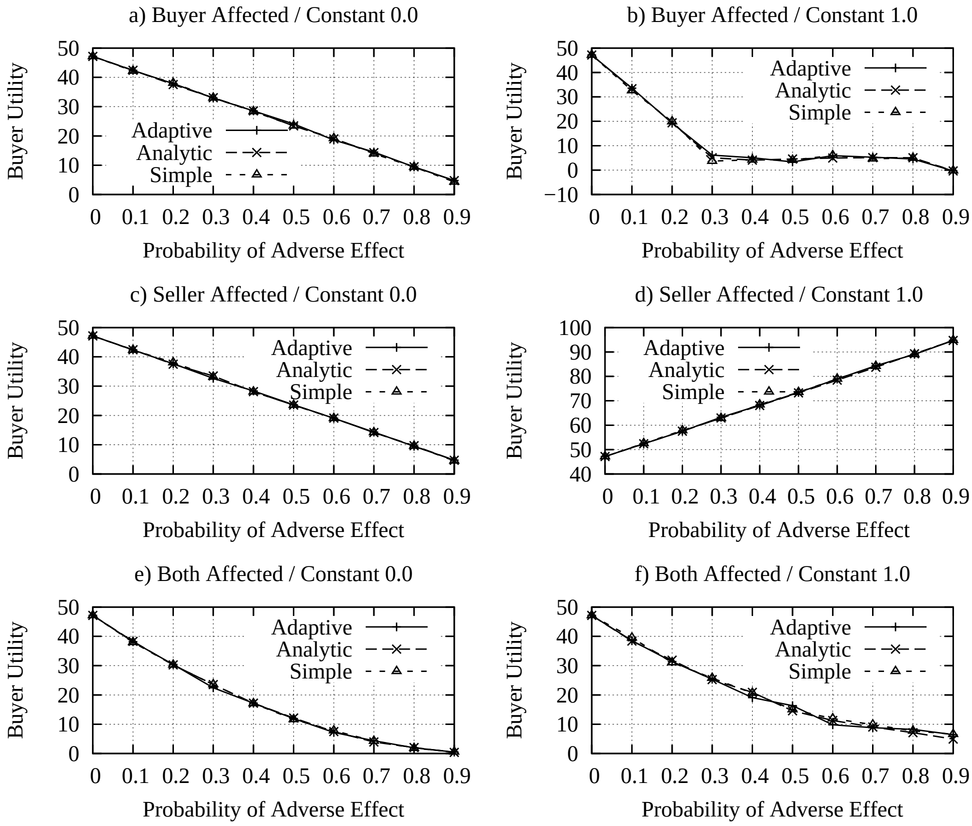

FIGURE 10.6: The performance of different concurrency strategies with the *Adaptive Counter* tactic and the *Expected Utility* opponent selection strategy (no future offers) (Hypothesis 22).

We will now move to discussing differences between the two advanced concurrency strategies (*Adaptive* and *Analytic*). The results of statistical testing for the case where the *Adaptive Counter* tactic is used at the `Negotiator` level are shown in table 10.2. The more plusses there are the more likely it is *Adaptive* beats the *Analytic* concurrency strategy in that specific setting. It is clear that, here, the *Adaptive* concurrency strategy beats the *Analytic* in most cases when the *Random* or *Quality* opponent selection is used, but not with the *Expected Utility.* This is consistent with hypothesis 22.c. We have therefore the results needed to accept the whole hypothesis 22.

We will now discuss the cases with the *Random Counter* tactic. Here, the performance of different concurrency strategies in different settings are shown in the figures 10.7–10.9 and the results of the statistical tests are summarised in table 10.3. As expected, the results are very different from the cases where *Adaptive Counter* tactic was used.

Affected Parties	Constant 0.0										Constant 1.0									
	Effect Probability										Effect Probability									
	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
Random Opponent Selection																				
Buyer	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	++	-	-	−−−	−−−	-	-	-	−	−−−
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	−−−	-	-	-	+++	+++	−−−
Seller	+++	+++	+++	+++	+++	+++	+++	++	+++	+++	++	-	-	-	-	-	-	-	-	-
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++
Both	+++	+++	+++	+++	++	++	+++	+++	-	-	++	-	-	++	-	-	-	−	-	-
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	-	+++	-	++
Quality Opponent Selection																				
Buyer	+++	+++	-	-	-	-	-	-	-	+	-	-	-	-	-	−−	-	-	-	−
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	-	-	−	++	-	+++	-
Seller	+++	+	+	++	+	-	+	-	-	-	-	-	-	-	-	-	-	-	-	-
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++
Both	+++	+	++	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
	+++	+++	+++	+++	+++	+++	+++	+++	+	-	+++	+++	+++	+++	+++	++	-	-	-	+
Expected Utility Opponent Selection																				
Buyer	-	-	-	-	+	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
	-	-	-	-	-	+	-	-	-	-	-	-	-	+	-	-	-	-	-	-
Seller	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
Both	-	-	-	-	-	-	-	−	-	-	-	-	-	-	-	-	-	-	-	-
	-	-	-	-	-	-	−−	-	-	-	-	-	-	-	-	-	-	-	-	-

Table 10.1: The performances of the advanced concurrency strategies as compared to the best *Simple* concurrency strategy (the *Adaptive Counter* tactic without future offers) (Hypothesis 22.a–b). The *Analytic* vs. the Best of the *Simple* on the top row and the *Adaptive* vs. the Best of the *Simple* on the bottom row in each case.

Affected Parties	Constant 0.0										Constant 1.0									
	Effect Probability										Effect Probability									
	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
Random Opponent Selection																				
Buyer	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	-	+++	-	-	++	+++	-
Seller	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++
Both	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	-	-
Quality Opponent Selection																				
Buyer	+++	+++	+++	+++	+++	+++	+++	+++	+++	-	+++	+++	+++	-	-	-	++	-	+++	+
Seller	+++	+++	+++	+++	+++	+++	+++	+++	+++	++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+
Both	+++	+++	+++	+++	+++	+++	+++	+++	+	-	+++	+++	+++	+++	+++	+++	-	-	-	++
Expected Utility Opponent Selection																				
Buyer	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
Seller	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
Both	-	-	-	−−	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-

Table 10.2: The relative performances of the *Adaptive* and *Analytic* concurrency strategies (the *Adaptive Counter* tactic without future offers) (Hypothesis 22.c).

In more detail, we go through the results row by row again, so we start with the *Analytic* concurrency strategy. It is able to beat the *Simple* strategies in most situations when the *Random* opponent selection and the *Constant 0.0* fee are used and also in most cases with the *Constant 1.0* policy when the seller is affected. However, it loses to the *Simple* strategies in many cases starting from those where the *Constant 1.0* policy is used and the buyer is affected (cases 0.3, 0.4 and 0.9 when *Random* and $0.3 - 0.5$ and 0.9 when the *Quality* opponent selection is used). In the cases with a very high buyer effect probability (0.9), the explanation follows from the inaccurate estimates for the contract times after the last matching and the fact that the *Random Counter* tactic is stuck with some opponents using the *Random* tactic also with *Expected Utility* opponent selection. This means that the performance of the *Simple 1* strategy deteriorates to the same level with the advanced strategies, not like with the *Adaptive Counter* where the advanced strategies improved their performance. This is because the buyers here often encounter sellers using the *Random* tactic.

However, here (unlike before) the *Analytic* strategy gets thoroughly beaten also in the cases where the seller or both parties are affected, the *Constant 0.0* and the *Quality* opponent selection is used. It seems that this effect also translates to the cases where the buyer is affected, at least when the probability of effect is very small. This last result is somewhat peculiar at first because we expected the *Simple* strategies to beat the advanced strategies where the contracts are formed earlier and this has an impact on the expected utility. However, there is no such effect when the buyer is affected and the *Constant 0.0* fee is used because it is free to exit the contracts. We did mention that there might be a weaker advantage for getting into contracts later and that was because a large number of very good quality providers are entering during the run, so when the *Simple 1* often fails once or twice, it gets access to these opponents with higher qualities. The difference in performance is not huge, but because the variations in these cases are quite small, even this little can be enough. When the variation increases with the effect probabilities this small advantage drowns mostly in the noise. All these results are consistent with hypothesis 23.a.

The *Adaptive* concurrency strategy is unable to do much better than the *Analytic* one in many situations. As with the *Analytic*, the *Adaptive* strategy does very well when the *Random* opponent selection and the *Constant 0.0* decommitment policy are used and it does do slightly better than the *Analytic* when the *Constant 1.0* decommitment policy is used: it consistently beats the competition when the seller is affected and when the buyer is affected with a very low probability or when both are affected with low to medium probability. However, it is beaten by the best *Simple* strategies in many cases where the *Constant 1.0* policy is used both when the buyer is affected (0.3, 0.4 and 0.9 when the *Random* and with 0.3, 0.5 and 0.9 when the *Quality* opponent selection is used) but interestingly, also when both are very likely to be affected ($0.7-0.9$ when the *Random* and 0.9 when the *Quality* opponent selection is used). And in addition, of course, it will get a thorough beating when the seller or both are affected and the *Quality* opponent selection and the *Constant 0.0* is used. The explanations are analogous to the case with the *Analytic* strategy and will not be repeated here. All these results are consistent with hypothesis 23.b.

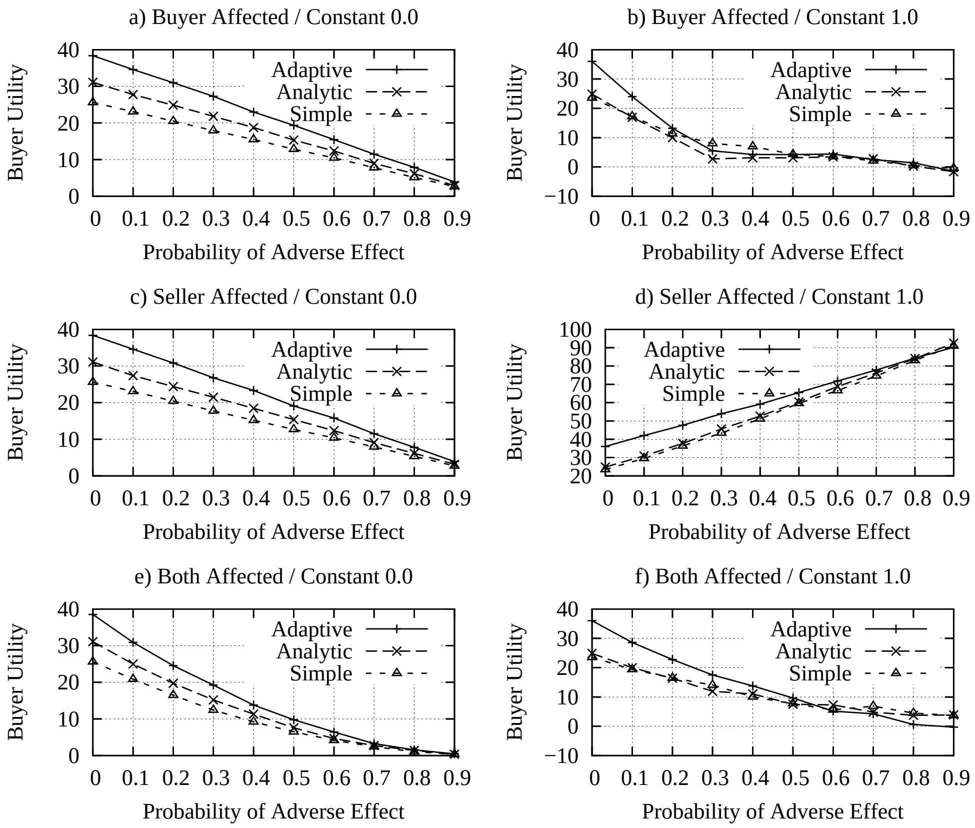

FIGURE 10.7: The performance of different concurrency strategies with the *Random Counter* tactic and the *Random* opponent selection strategy (Hypothesis 23).

We will now move to discussing the differences between the two advanced concurrency strategies (*Adaptive* and *Analytic*). The results of statistical testing for the case where the *Random Counter* tactic is used at the `Negotiator` level are shown in table 10.4. The more plusses, the more probably *Adaptive* beats the *Analytic* concurrency strategy in the specified setting. From the results, it is clear that, here, the *Adaptive* concurrency strategy beats the *Analytic* only when the *Random* opponent selection is used but not much else where. Even in the *Random* opponent selection, it has trouble when the *Constant 1.0* decommitment policy is used and the effect probabilities are high. It even gets beaten occasionally there in the extreme cases. The losses can be explained by the fact that the *Adaptive* strategy can find contracts earlier than the *Analytic* strategy. With the *Quality* or *Expected Utility* opponent selection, there is very little to go between the two strategies. The seller providing the highest expected utility is going to be the

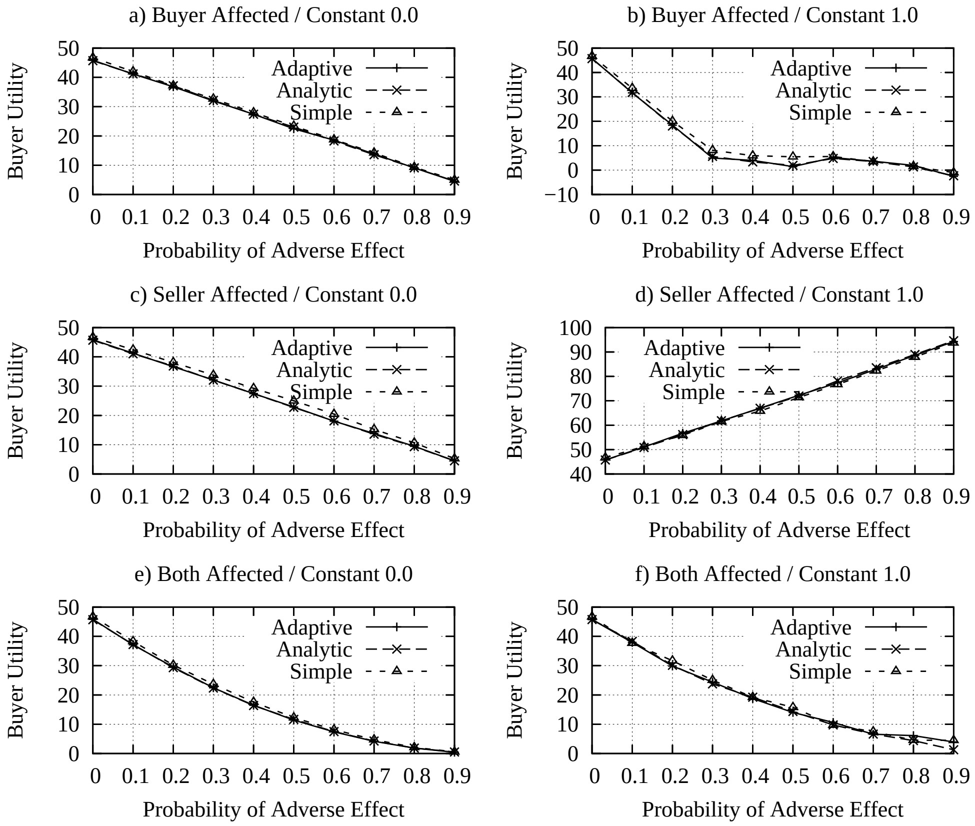

FIGURE 10.8: The performance of different concurrency strategies with the *Random Counter* tactic and the *Quality* opponent selection strategy (Hypothesis 23).

first opponent in the negotiation queue that any result can be achieved, so both strategies will be negotiating mostly with the same opponents.

However, with free decommitments and no adverse effects, the *Adaptive* strategy seems to enjoy an edge. This is because the *Adaptive* strategy is able to use the existence of other providers offering good deals to sometimes increase the expected utility it can get from the contracts above that of the best single negotiation. With positive fees, such a strategy is riskier. The effect also requires the best opponents using the *Random* negotiation tactic and having very similar qualities. Moreover, the possibility of adverse effects may make the effect weaker and that is why we only see the difference when the *Constant 0.0* decommitment policy is used and there are no effects. The only exception involves 0.1 seller effect probability when the *Expected Utility* opponent selection is used. All this is consistent with hypothesis 23.c and we can therefore accept the whole of hypothesis 23.

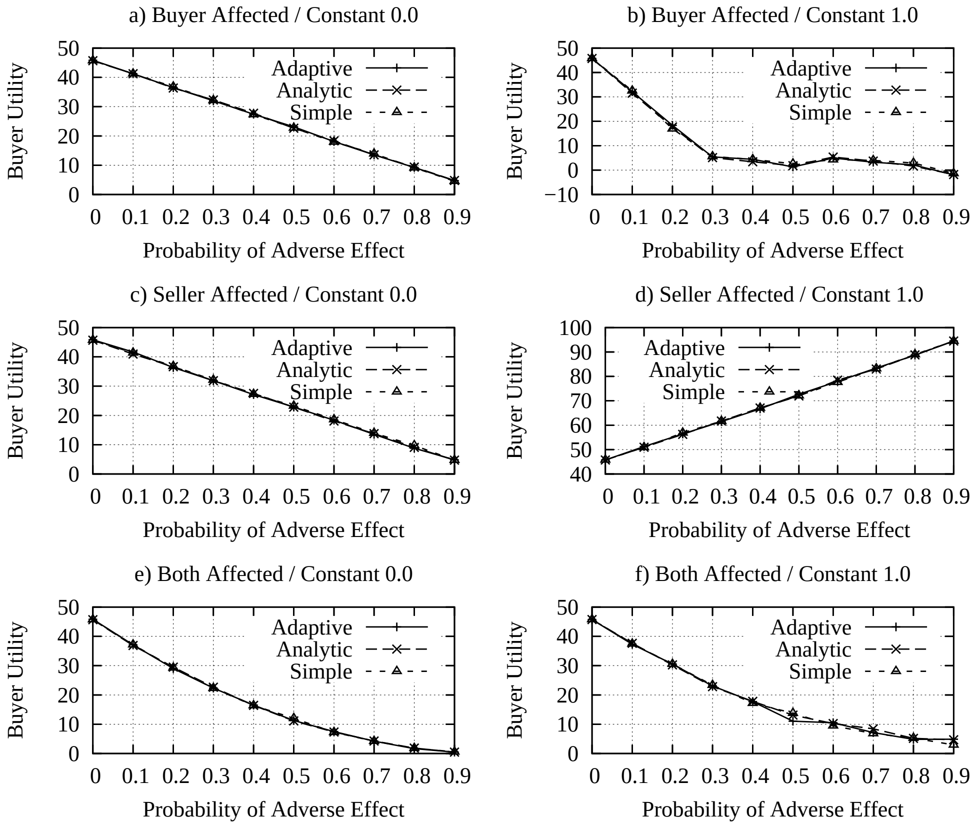

FIGURE 10.9: The performance of different concurrency strategies with the *Random Counter* tactic and the *Expected Utility* opponent selection strategy (Hypothesis 23).

10.2.3.2 Full Tactic Information with Future Offers

We will now move to cases where the buyer takes the offers it might receive later into account when deciding whether or not to take a given contract. As we discussed in section 10.2.1.2, this does change the nature of the exercise from the situation we discussed earlier (no future offers). Here, the buyer may turn down contracts that seem beneficial (yield positive expected utility), because it considers it likely that a better offer will be forthcoming in a later negotiation. As explained, this will shift the focus to what sort of opponents will the buyer be able to encounter when it matters, i.e. when the most beneficial contracts are made.

As in the case without the future offers, we will discuss here how advanced strategies fare against the *Simple* ones (hypothesis 24) and the relative performances of the advanced strategies (hypothesis 25). In addition, we will discuss the effect of opponent selection (hypothesis 26) and future offers (hypothesis 27) and

Affected Parties	Constant 0.0										Constant 1.0									
	Effect Probability										Effect Probability									
	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
Random Opponent Selection																				
Buyer	+++	+++	+++	+++	+++	+++	+++	+++	+++	+	-	-	-	−−−	−−−	-	-	-	-	−−−
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	-	−	−−	-	-	-	+	−−
Seller	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	-	-	-	++	+	-	+++	++	+	+++
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+	-
Both	+++	+++	+++	+++	+++	+++	++	+	+++	-	-	-	-	−	-	-	-	−	-	-
	+++	+++	+++	+++	+++	+++	+++	+++	+++	-	+++	+++	+++	+++	+++	+	-	−−	−−−	−−−
Quality Opponent Selection																				
Buyer	−−−	−−−	-	-	-	-	-	−	-	-	−−−	−−	−	−−	−−	−−−	-	-	-	−−−
	−−−	−−−	-	−	-	−−	-	-	-	−	−−−	−−	−	−−	−	−−−	-	-	-	−−−
Seller	−−−	−−−	−−−	−−−	−−−	−−−	−−−	−−−	−−−	−−−	−−−	-	-	-	++	-	+++	+++	+++	+++
	−−−	−−−	−−−	−−−	−−−	−−−	−−−	−−−	−−−	−−−	−−−	++	-	-	++	+	-	+	+	+
Both	−−−	−−−	−−	−−−	−−−	−	−−	−	-	-	-	−−−	-	-	-	-	-	-	-	−
	−−−	−−−	−	−−−	−−	−−	−−	−−	-	-	-	−−−	-	−	-	-	-	-	-	-
Expected Utility Opponent Selection																				
Buyer	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
	+++	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
Seller	-	−	-	-	-	-	-	-	−	-	-	-	-	-	-	-	-	-	-	-
	+++	-	-	-	-	-	-	-	−−−	-	-	-	-	-	-	-	-	-	-	-
Both	-	-	-	-	-	−−	-	-	-	-	-	-	-	-	-	-	-	-	-	-
	+++	-	-	-	-	−	-	-	-	-	-	-	-	-	-	-	-	-	-	-

Table 10.3: The performances of the advanced concurrency strategies as compared to the best *Simple* concurrency strategy (the *Random Counter* tactic without future offers) (Hypothesis 23.a–b). The *Analytic* vs. the Best of the *Simple* on the top row and the *Adaptive* vs. the Best of the *Simple* on the bottom row in each case.

Affected Parties	Constant 0.0										Constant 1.0									
	Effect Probability										Effect Probability									
	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
Random Opponent Selection																				
Buyer	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	++	-	-	-	-	++	-
Seller	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	-	-
Both	+++	+++	+++	+++	+++	+++	+++	+++	-	-	+++	+++	+++	+++	+	-	-	-	−	−−
Quality Opponent Selection																				
Buyer	+++	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
Seller	+++	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	−	-	-	-
Both	+++	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	+
Expected Utility Opponent Selection																				
Buyer	+++	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	++	-
Seller	+++	+++	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
Both	+++	-	-	-	-	-	-	-	-	-	-	-	-	-	-	++	-	-	-	-

Table 10.4: The relative performances of the *Adaptive* and *Analytic* concurrency strategies (the *Random Counter* tactic without future offers) (Hypothesis 23.c).

we will also investigate the relative performances of the two negotiation tactics (hypothesis 28).

We start by comparing the performances of the advanced and the *Simple* concurrency strategies. The results for cases where the *Random Counter* tactic is used at the `Negotiator` level are in figures 10.10–10.12 and in table 10.5. From this table, it is clear that the advantage of the advanced concurrency strategies has completely different patterns here than in the case with no future offers. The advantage is strong with both the *Quality* and the *Expected Utility* opponent selection and actually it is in many cases somewhat stronger than with the *Random* opponent selection where the advantage was the strongest earlier. This is because with the future offers considered, most contracts are entered into later during the experiment and the advanced strategies will be able to benefit more from their willingness to negotiate only when good contracts are to be found. This means that with more advanced opponent selection strategies, the *Simple* strategies will have negotiated with many of the best opponents when the contracts are formed, whereas the advanced strategies are able to pick the best opponents at that point. With the *Random* opponent selection, the difference is often less pronounced, because the *Simple* strategies negotiate too early with very bad and mediocre opponents as well as some good ones and also when the time comes, the difference in what the different negotiation strategies can actually see is smaller.

There are a few deviations from this general trend. The first are the cases where the *Constant 0.0* decommitment policy is used and only the buyer can be affected. Here, the advantage of the advanced strategies is often sporadic at best. This is because when the fee is zero and only the buyer is affected, the contract time does not matter (adverse impact and no contract produce the same outcome, zero utility) and, therefore, the advanced strategies are usually unable to beat the best *Simple* strategies. There are even some cases with the *Random* opponent selection, where the *Simple* strategies outperform the advanced strategies. This is because the advanced strategies will encounter fewer opponents than the best *Simple* strategies and may, therefore, be unable to find as good opponents as the *Simple* strategies. However, when the probability of adverse effect increases, the best contracts are usually formed nearer the deadline and this means that the time window where almost all contracts are formed is small and, therefore, it matters not what has happened before.

The other deviation is in the cases with the *Constant 1.0* decommitment policy, when only the seller is affected. The advantage seems to show up only with medium to high effect probabilities and it vanishes altogether when the *Expected Utility* opponent selection is used. This is quite understandable. In these cases, the non-performance is actually preferred to performance and, therefore, the most important goal for the buyer is get into a contract and preferrably as soon as possible. This means that with the *Random* opponent selection, where many of the encountered opponents are bad or mediocre and therefore eliminated when future offers are considered, the best *Simple* strategy is typically to have a large number of negotiations to maximise the chances of getting into contracts quickly. However, this may mean an occasional extra contract and when the fee is high, this can be risky. On the other hand, with the *Quality* opponent selection, the best approach is to use very few negotiations because the *Quality* opponent selection means that there are many more good opponents to meet. This means, however, problems with getting into contracts because the small number of negotiations may mean delays. In both cases, the push to get into contracts quickly is less with low probabilities of seller effect and, therefore, also the differences are often smaller. This means especially that the advanced strategies will not necessarily find it to be in their interest to get into contracts in the first matching. The situation will change when the probability increases. When the *Expected Utility* opponent selection is used, in contrast, the best opponent is the first opponent for all strategies and both advanced and the *Simple 1* strategies will enter into a contract around the same time. The type of opponents they encounter later is similar and their choices are too. Especially with mid-to-high seller effect probabilities, they will both negotiate only with the first available opponent.

The third case is with *Constant 1.0* when both parties can be affected. Here, the advantage of the advanced strategies over the *Simple* ones is quite weak with the *Random Counter* negotiation tactic, especially with the *Random* opponent selection. The explanation here is that when both are affected with the same probability and a *Constant* decommitment policy is used, the chance of huge fees through adverse effects cancel each other out. On expectation, it is equally like to get or having to pay the fee, so it does not have any effect on the optimal strategy. The only relevant factor in this respect is that any non-performance means no utility (on expectation) and, therefore, minimising the probability of failure might always be a good option. However, with the *Random* opponent selection, this gives the advanced strategies no significant edge. The only risk associate with the high decommitment fee is the risk of extra contracts, but with the *Random Counter* even that risk is usually relatively low and especially with future offers most bad opponents are automatically removed until the last matching. This means that the *Simple* strategies can do well even with relatively high number of negotiations which in turn means that it will be able to find relatively good contracts also when all contracts are formed near the deadline.

When the opponent selection improves, the advanced strategies improve their performance relative to the best *Simple* strategies. This is because the *Simple* strategies waste too many good opponents by negotiating with them too early and/or they will have trouble ensuring a contract with limited number of negotiations after the last matching because of the uncertainties involved in negotiations after the last matching.[29] Also the probabilities of success can be relatively low with

[29] These include earlier deadlines that are also more difficult to guess.

Affected Parties	Constant 0.0										Constant 1.0									
	Effect Probability										Effect Probability									
	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
Random Opponent Selection																				
Buyer	–	-	––	-	-	-	-	-	-	-	+	-	-	+	-	+++	++	-	-	–––
	-	-	–	-	-	-	-	-	-	-	++	+++	+++	+++	+++	+++	+++	+++	++	––
Seller	–	-	-	-	-	+	-	+++	+++	+++	+	-	-	-	-	+++	-	+	++	+++
	-	-	-	+++	+++	+++	+++	+++	+++	+++	++	+	+	-	+	++	++	+	+++	+++
Both	–	-	-	-	-	+	++	+	-	+	+	-	-	-	-	+	-	-	-	-
	-	-	-	+	+	++	+++	+++	+++	+++	++	-	-	-	-	-	-	-	-	-
Quality Opponent Selection																				
Buyer	++	-	-	-	-	-	-	-	-	-	+++	+++	+++	+++	+++	+++	+++	+++	-	––
	+++	++	++	-	-	-	-	-	-	-	+++	+++	+++	+++	+++	+++	+++	+++	-	–
Seller	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	-	-	-	+	++	+	+++	++	+
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	-	-	-	-	+	++	+++	+	-
Both	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	++	+++	+++	+++	+++	+++	-	-	-
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	++	++	+++	++	++	+++	-	-	-
Expected Utility Opponent Selection																				
Buyer	++	-	-	-	-	-	-	-	-	-	+++	+++	+++	+	+	-	-	-	-	-
	+++	+	+	+	-	+	-	-	-	-	+++	+++	++	++	++	-	-	-	-	-
Seller	++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	-	-	-	-	-	-	-	-	-
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	-	-	-	-	-	-	-	-	-
Both	++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	-	++	++	+	+++	+	-
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	++	+	+++	++	+	+++	++	-

TABLE 10.5: The performances of the advanced concurrency strategies as compared to the best *Simple* concurrency strategy (the *Random Counter* tactic with future offers) (Hypothesis 24). The *Analytic* vs. the Best of the *Simple* on the top row and the *Adaptive* vs. the Best of the *Simple* on the bottom row in each case.

very high effect probabilities even near the deadline[30] and the variance is often quite high (because of the high decommitment fee).

And of course there is also the case where the buyer is affected, the *Constant 1.0* policy is used and the effect probability is very high (0.9). Here, the advanced strategies lose to the *Simple 1* strategy with the *Random* opponent selection and to a lesser extent with the *Quality* strategies. However, there is no difference with the *Expected Utility* opponent selection for the very simple reason that the best opponent is always the first one in the negotiation queue and negotiating with anybody else is never useful. However, with the less advanced opponent selection strategies, the advanced strategies' ability to find these problematic sellers (whereas the *Simple 1* often does not meet them) can still present a problem just as it did in the cases with no future offers. This was to be expected because all the contracts are formed at the last minute (generally, after the last matching) and there are no future offers in that last matching.

Now, the results with the *Adaptive Counter* tactic, shown in figures 10.13–10.15 and table 10.6, are quite similar to the results we just discussed. However, there are some interesting differences. First of all, we get the same trends and exceptions we discussed earlier. The advantage is stronger with more advanced opponent selection strategies, because, as before, the *Simple* strategies negotiate with many

[30]For example, when both parties have an adverse effect with probability of 0.5, only 25% of the contracts are going to be performed.

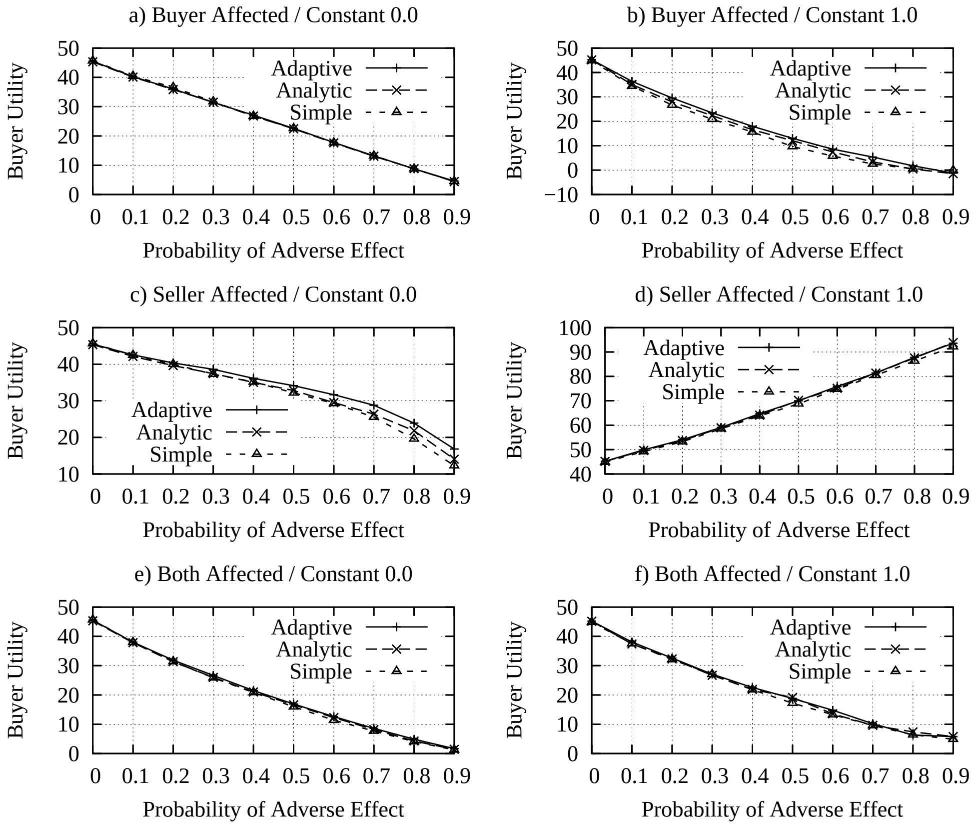

FIGURE 10.10: The performance of different concurrency strategies with the *Random Counter* tactic, the *Random* opponent selection strategy and future offers (Hypotheses 24–25).

good opponents too early, whereas the advanced strategies wait for their opportunities. This means that they can set their thresholds higher. With the *Adaptive Counter* tactic, the differences are less pronounced, the differences are very strong (where they exist) in most cases already with the *Random* opponent selection. Especially the *Adaptive* concurrency strategy does quite well even with the *Constant 0.0* decommitment policy. This is because the *Adaptive* strategy is so much more efficient in picking up the best opponents than the *Simple* (or *Analytic*) strategies that this more than compensates for encountering fewer opponents. The *Analytic* strategy does improve its position too, since it only loses to the *Simple* strategies (due to meeting less opponents), when no adverse effects occur.

As for the exceptions, we can conclude from table 10.6 that when the fee is zero and only the buyer is affected, we have only sporadic advantage for the advanced strategies. Here, however, the advantage is somewhat stronger with both the *Quality* and the *Expected Utility* opponent selection than it was with the *Random Counter* negotiation tactic. This is because the *Adaptive Counter* tactic is

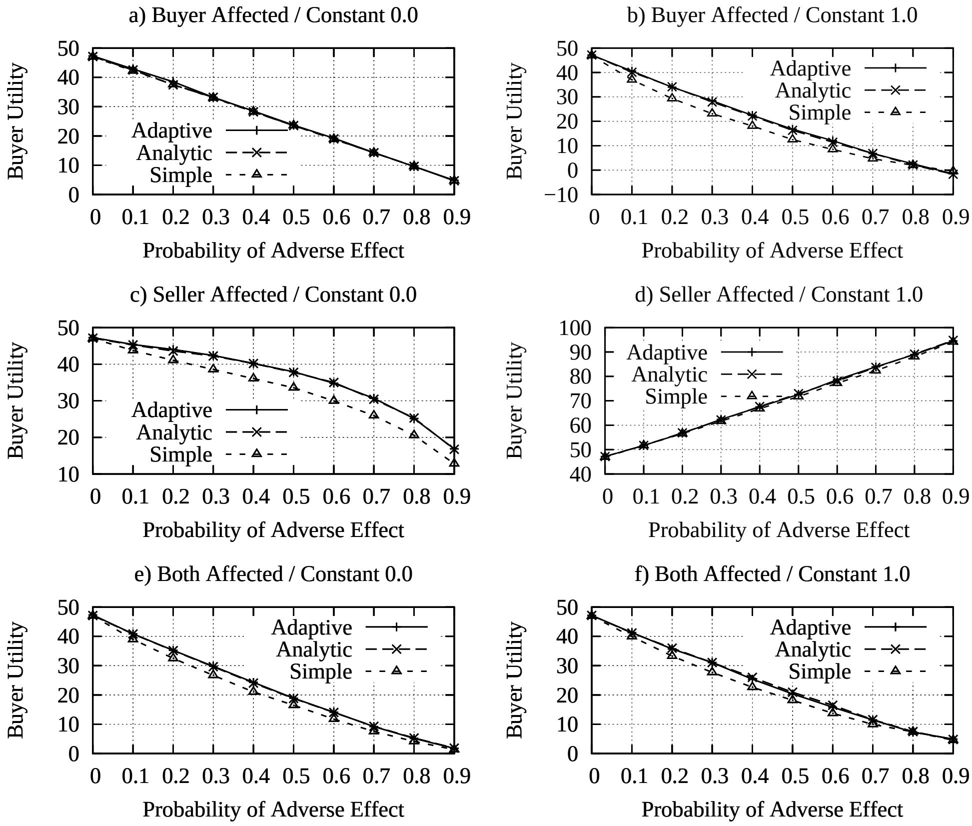

FIGURE 10.11: The performance of different concurrency strategies with the *Random Counter* tactic, the *Quality* opponent selection strategy and future offers (Hypotheses 24–25).

holding on to the good opponents and, even more importantly, many of these good opponents will be using the *Exponential Time-Dependent* tactic, which the *Adaptive Counter* tactic can exploit. The differences between excellent and merely very good are therefore greater. So even if waiting does not improve the buyer's chances of success, it can mean finding better opponents to negotiate with. This allows some advantage to the advanced strategies.[31]

The other effect, which is also stronger, is that when the fee is high and only the seller can be affected, the advantage of advanced strategies dissipates with more sophisticated opponent selection. The advantage is still clear in all cases with the *Random* opponent selection, however there are cases with no performance difference in the *Quality* opponent selection and there is no advantage whatsoever with the *Expected Utility* opponent selection. As before, this is because it becomes easier for the *Simple* strategies to find a good opponent to quickly get

[31]Using the future offers will mean that getting into a contract is delayed also when the buyer is affected. This had less effect with the *Random Counter* tactic because all high quality providers using a non-behavioural tactic will produce a roughly similar result.

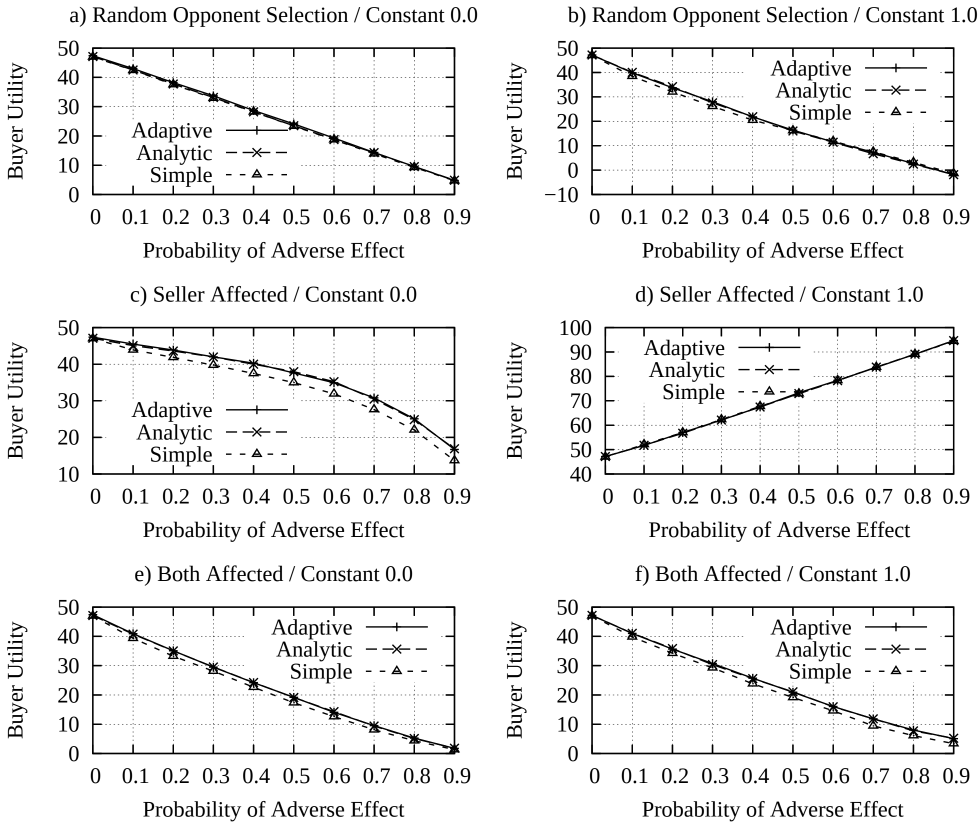

FIGURE 10.12: The performance of different concurrency strategies with the *Random Counter* tactic, the *Expected Utility* opponent selection strategy and future offers (Hypotheses 24–25).

into a contract with. With the *Adaptive Counter* negotiation tactic, however, the advanced strategies can hold their advantage longer, because even with a very similar quality, the tactic's adaptation means that outcomes will be very different with different opponent tactics. Therefore choosing only the good opponents can be important. However, as before, with the *Expected Utility* opponent selection, any advantage is lost because all parties negotiate with the best opponent and it alone.

Earlier, the third exception was the cases with *Constant 1.0* when both parties can be affected. With the *Adaptive Counter* negotiation tactic this does not become an issue at any point, because the risk of an extra contract is very high and the *Simple* strategies with many negotiations cannot perform well as a result. This means that the advanced strategies will be able to outperform the *Simple* strategies also with the *Random* opponent selection. However, as earlier, the *Adaptive Counter* negotiation tactic is unable to avoid too many decommitments with very high buyer effect probabilities when the *Constant 1.0* decommitment policy is used.

Thus, is no difference in this respect with the *Expected Utility* opponent selection strategy. This and all the other findings discussed are consistent with hypothesis 24 and we can therefore accept it.

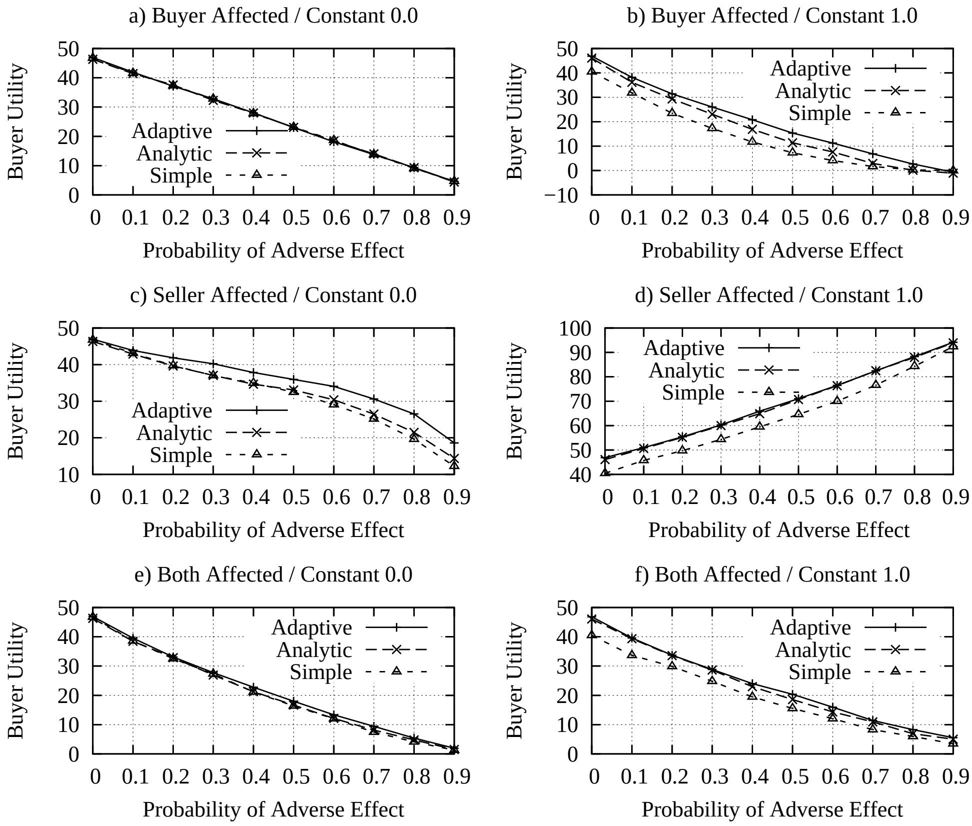

FIGURE 10.13: The performance of different concurrency strategies with the *Adaptive Counter* tactic, the *Random* opponent selection strategy and future offers (Hypotheses 24–25).

The next step is to consider the differences between the *Adaptive* and *Analytic* concurrency strategies. Many of the differences we had earlier are removed, because the future offers remove most bad and mediocre opponents. However, they do not do this perfectly because they operate on averages and there may therefore be more than one opponent that exceeds this threshold in an actual case. Also it seems that in many situations, the highest reasonable thresholds that minimise the number of extra opponents is not always the best idea, because it means that in some cases there may be no opponents at all to exceed the threshold. Often, especially for the *Analytic* strategy, lowering the thresholds somewhat improves the performance. Moreover, in the last matching, the threshold is zero, so if the buyer has no contract before then, the *Adaptive* strategy may be able to get its chance to shine, because the *Adaptive* concurrency strategy still usually has an edge over

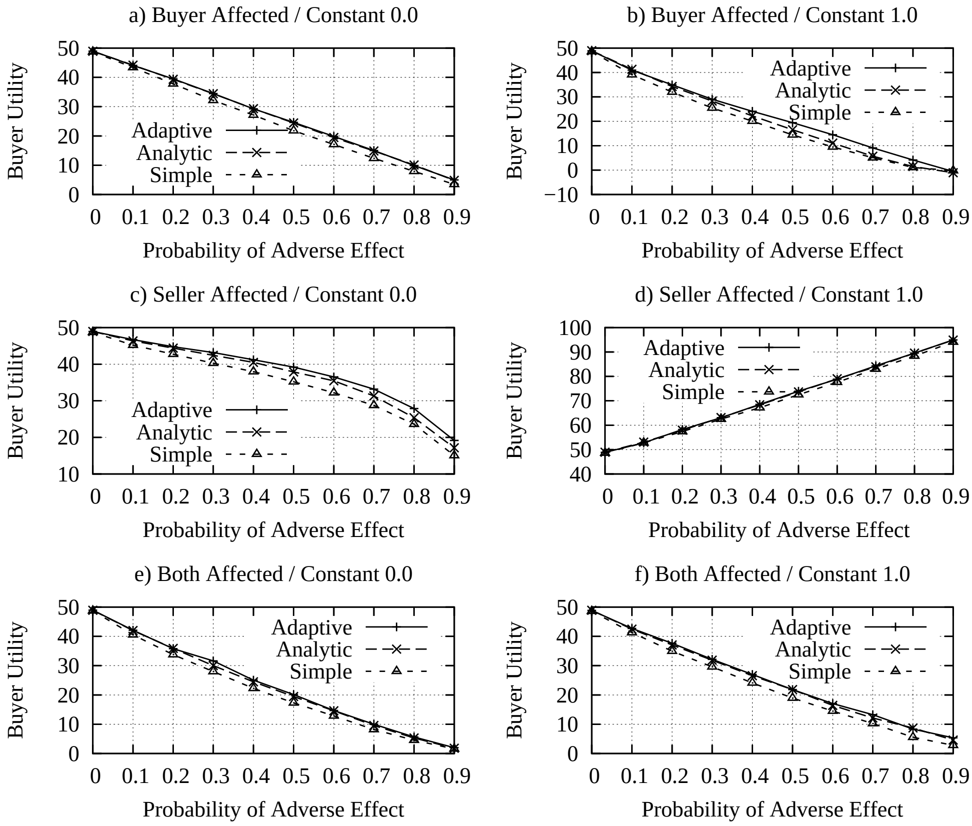

FIGURE 10.14: The performance of different concurrency strategies with the *Adaptive Counter* tactic, the *Quality* opponent selection strategy and future offers (Hypotheses 24–25).

the *Analytic* one when there are multiple opponents (although the advantage is not that great always).[32]

In more detail, table 10.7 shows the relative performance levels of the advanced strategies with the *Random Counter* tactic. There seems to be a consistent difference only with the *Random* opponent selection and when one of the parties is affected (seller when fee is zero and buyer when fee is one); that is, when it is clearly significantly in the buyer's interest to delay. It seems that this is where the difference is most clearly in the *Adaptive* strategy's favour. There is also some of this happening when both are affected (*Constant 0.0*), but otherwise it is just single instances. However, the case with the *Constant 0.0* decommitment fee and very low effect probabilities is present as it was without future offers and the explanation is the same: when the probability of adverse effect is near zero, the *Adaptive* concurrency strategy will be able to use negotiations with multiple

[32]Sometimes the difference between the two is relatively small and sometimes it may not be in the buyer's interest to negotiate with more than one opponent.

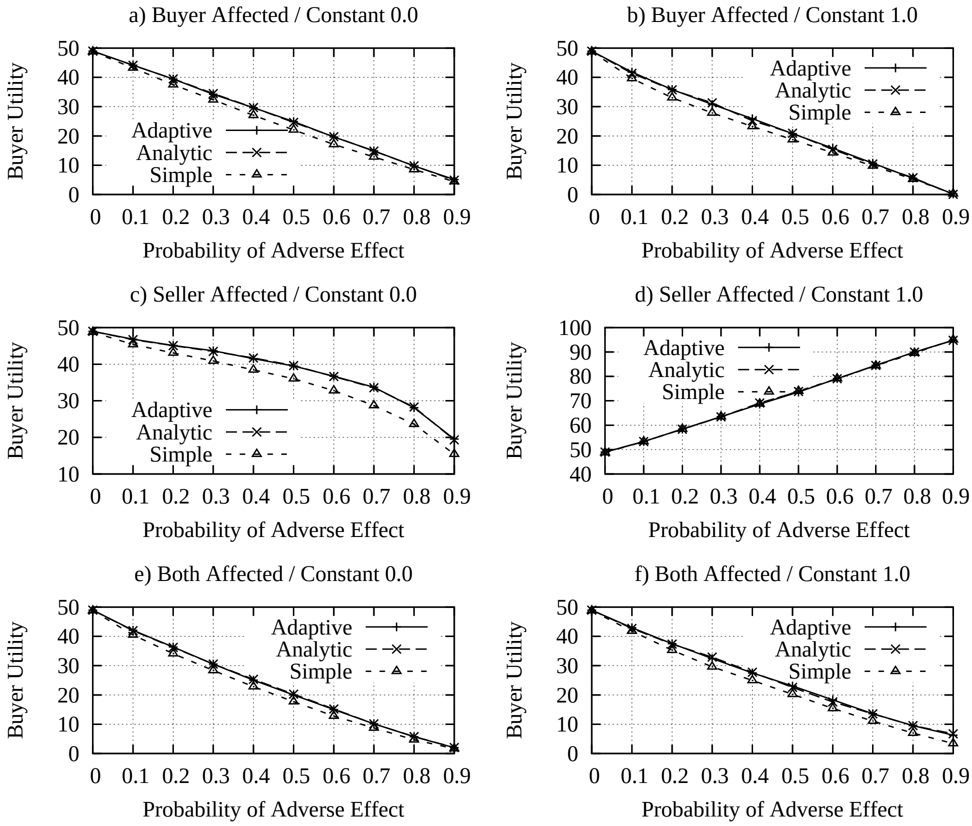

FIGURE 10.15: The performance of different concurrency strategies with the *Adaptive Counter* tactic, the *Expected Utility* opponent selection strategy and future offers (Hypotheses 24–25).

sellers to increase its expected utility. With higher effect probabilities this will not be possible, because the effect decreases the usefulness of this trick by lowering the improvement in expected utility. Here, however, because the future offers often make the buyer wait, this effect is lower (because the effect probabilities are lower when the contracts are made nearer the deadline). Therefore, we get a couple of extra signs of difference in other cases.

In this vein, table 10.8 shows the performance differences between the *Adaptive* and *Analytic* strategies for the case where the `Negotiator` uses the *Adaptive Counter* tactic. As before, we get slightly stronger differences here than we did with the *Random Counter* tactic.[33] Specifically, there is a clear difference in the two cases where we had the difference also with the *Random Counter* tactic, namely, seller affected (*Constant 0.0*) and buyer affected (*Constant 1.0*). In addition, we have a clear difference when both are affected (*Constant 0.0*). Moreover, there are

[33] As before, this difference between the negotiation tactics exists because the *Adaptive Counter* tactic adapts too all opponent tactics and especially important is that it will be able to exploit the *Exponential Time-Dependent* tactic unlike *Random Counter*.

Affected Parties	Constant 0.0										Constant 1.0									
	Effect Probability										Effect Probability									
	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
Random Opponent Selection																				
Buyer	– – –	-	-	-	-	-	-	-	-	-	+++	+++	+++	+++	+++	+++	+++	++	-	– – –
	++	-	-	-	-	-	-	-	-	-	+++	+++	+++	+++	+++	+++	+++	+++	+++	–
Seller	– – –	-	-	-	-	-	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++
	++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++
Both	– – –	-	-	-	-	-	-	++	+++	+	+++	+++	+++	+++	+	+	-	++	-	-
	++	+++	-	-	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	++	+++	+++	++	+	++
Quality Opponent Selection																				
Buyer	++	+++	-	-	-	-	-	++	+++	+++	+++	+++	+++	+++	++	+++	++	-	-	–
	+++	++	-	-	-	-	+	+++	+++	+++	+++	++	+++	+++	+++	+++	+++	+++	+++	-
Seller	++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	-	-	-	++	+	++	+	++	++
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	-	+	-	+++	+++	++	+++	+++	++
Both	++	+++	+++	+++	+++	+++	+++	+++	+++	-	+++	-	+	++	++	++	+	+	++	++
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+	+++	++	++	++	++	+++	++	+++
Expected Utility Opponent Selection																				
Buyer	+++	-	+	+++	++	+	+	-	-	-	+++	+++	+++	+++	+++	+++	+	-	-	-
	+++	-	-	++	++	++	+	+	-	-	+++	+++	+++	+++	+++	+++	++	+	-	-
Seller	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	-	-	-	-	-	-	-	-	-
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	-	-	-	-	-	-	-	-	-
Both	++	+++	+++	+++	+++	+++	+++	+++	+++	++	+++	+	+++	+++	+++	+++	+++	+++	++	-
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+	+++	++	+++	+++	+++	+++	+++	+++	++	-

TABLE 10.6: The performances of the advanced concurrency strategies as compared to the best *Simple* concurrency strategy (the *Adaptive Counter* tactic with future offers) (Hypothesis 24). The *Analytic* vs. the Best of the *Simple* on the top row and the *Adaptive* vs. the Best of the *Simple* on the bottom row in each case.

Affected Parties	Constant 0.0										Constant 1.0									
	Effect Probability										Effect Probability									
	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
Random Opponent Selection																				
Buyer	++	-	-	-	-	-	-	-	-	-	-	+++	-	+	++	+	+	+++	+++	-
Seller	++	+	+++	+++	+++	+++	+++	+++	+++	+++	-	-	-	-	-	-	-	-	-	-
Both	++	-	+	+	-	-	-	-	+++	+	-	-	-	-	-	-	-	-	-	-
Quality Opponent Selection																				
Buyer	+++	-	+++	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
Seller	+++	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
Both	+++	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
Expected Utility Opponent Selection																				
Buyer	+++	+	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
Seller	+++	+	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
Both	+++	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-

TABLE 10.7: The relative performances of the *Adaptive* and *Analytic* concurrency strategies (the *Random Counter* tactic with future offers) (Hypotheses 25).

similar patterns also when the *Quality* opponent selection is used, so where the seller can be affected (*Constant 0.0*) and the buyer can be affected (*Constant 1.0*). We also get more than one plus when both are affected (*Constant 0.0*). This is all because the differences between the expected utilities vary more here and therefore it is likely that there are more opponents that provide utilities that are over averages, simply because every opponent is a potential contract partner, whereas with *Random Counter* half of the opponents are not considered at all. This means that the *Adaptive* concurrency strategy has more opportunities to

Affected Parties	Constant 0.0										Constant 1.0									
	Effect Probability										Effect Probability									
	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
Random Opponent Selection																				
Buyer	+++	+	-	+	-	-	-	-	-	-	+++	++	++	++	+++	+++	+++	+++	+++	-
Seller	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	-	-	-	++	-	-	-	-	-
Both	+++	+++	-	++	+++	+++	+++	+++	++	+++	+++	-	-	-	-	++	-	-	-	-
Quality Opponent Selection																				
Buyer	-	-	-	-	-	-	-	-	-	-	-	-	-	-	++	+++	+++	+++	+++	-
Seller	-	-	-	+	+	+++	++	+++	+++	+++	-	-	-	-	-	-	-	-	-	-
Both	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
Expected Utility Opponent Selection																				
Buyer	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
Seller	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
Both	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-

TABLE 10.8: The relative performances of the *Adaptive* and *Analytic* concurrency strategies (the *Adaptive Counter* tactic with future offers) (Hypothesis 25).

shine. All these observations are consistent with hypothesis 25 and we therefore accept it.

The third point of our interest is the role of the opponent selection. The results for the three concurrency strategies when the *Random Counter* tactic is used at the `Negotiator` level are shown in table 10.9. From this, it is clear that switching from the *Random* to the *Quality* opponent selection improves the situation in most cases. There are only a couple of exceptions. First, there are the cases with a very high buyer effect probabilities, which can be explained by the very few contracts that occur in these cases. Even if the opponents improve they are only rarely good enough to warrant negotiation. Second, the best of the *Simple* strategies are unable to improve where both parties can be affected with medium to high effect probabilities. This is because the probability of simultaneous contracts is so small that when the *Random* opponent selection is used, the buyer is able to use quite a few negotiations without a significant risk of extra contracts and, therefore, it will usually be able to find a good opponent to enter into contract with. All the *Quality* opponent selection does is to make the *Simple* strategies with many negotiations less useful, because they waste too many good opponents. The advanced strategies are able to benefit from the change, of course, because they usually do not have to worry about either of these forces. They will negotiate only when it is useful and when they do negotiate, they can negotiate with up to ten opponents, so finding a suitable opponent is all but guaranteed.

In contrast, there is practically no improvement on the advanced strategies when moving from *Quality* to *Expected Utility* opponent selection. This is because the

Affected Parties	Constant 0.0										Constant 1.0									
	Effect Probability										Effect Probability									
	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
Quality vs. Random Opponent Selection																				
Buyer	+++ +++ +++	+++ +++ +++	+++ +++ +++	+++ +++ +++	+++ +++ +++	+++ +++ +++	++ ++ ++	+++ +++ +++	+++ +++ +++	- - -	+++ +++ +++	+++ +++ +++	+++ +++ +++	++ +++ +++	+++ +++ +++	+++ +++ +++	+++ +++ +++	+++ +++ ++	++ +++ -	- - -
Seller	+++ +++ +++	+++ +++ +++	+++ +++ +++	+++ +++ +++	++ +++ +++	+++ +++ +++	+ +++ +++	- +++ +++	++ +++ +++	+ +++ -	+++ +++ +++	+++ +++ +++	+++ +++ +++	+++ +++ +++	+++ +++ +++	+++ +++ +++	+++ +++ +++	+++ +++ +++	+++ +++ +++	+++ +++ +++
Both	+++ +++ +++	+++ +++ +++	++ +++ +++	+ +++ +++	- +++ +++	- +++ +++	- +++ +++	- +++ +	- +++ -	- +++ -	+++ +++ +++	+++ +++ +++	- +++ ++	- +++ +++	- +++ ++	- ++ +	- +++ -	- ++ -	- - -	- - -
Expected Utility vs. Quality Opponent Selection																				
Buyer	- - +++	- - -	- - -	- - -	- - -	- - -	- - -	- - -	- - -	- - -	+ - -	++ - -	+++ - -	+++ - -	+++ - -	+++ - -	+++ - -	+++ - -	+ - -	- - -
Seller	- - +++	- - -	++ - -	+++ - -	+++ - -	+++ - -	+++ - -	+++ - -	+++ - -	++ - -	- - -	- - -	- - -	- - -	+ - -	++ - -	++ - -	+++ - -	+ - -	- - -
Both	- - +++	- - -	+ - -	+++ - -	+++ - -	+ - -	++ - -	+ - -	- - -	- - -	+ - -	- - -	- - -	++ - -	+ - -	- - -	- - -	- - -	- - -	- - -

Table 10.9: Effect of Opponent Selection (Random Counter Tactic with Future Offers) (Hypothesis 26). Best of Simple on the top row, Analytic in the middle and Adaptive on the bottom row.

Quality opponent selection often means that the buyer already meets very reasonable opponents (this is because the *Quality* opponent selection already orders opponents in descending order of expected utility). The *Expected Utility* opponent selection only removes the opponents using the behavioural tactics from the negotiation queue and the advanced strategies can remove these opponents on their own. However, for the simple strategies the move is often useful, because they can move to the *Simple 1* strategy with the *Expected Utility* opponent selection whereas this would not be that good a strategy and therefore we get consistent improvements in many situations. The exception is when the buyer is affected and there is the *Constant 0.0* policy in use (and this is because the strategies in general will not be able to make any improvements here).

The results for the other case, where the `Negotiator` uses the *Adaptive Counter* tactic are given in table 10.10. Again, the switch from the *Random* to the *Quality* opponent selection improves the situation in most settings. There are some improvements also when moving from the *Quality* to the *Expected Utility* opponent selection. Here, we get consistent improvements for the *Analytic* concurrency strategy for cases where only the seller can be affected (the *Constant 0.0* decommitment policy) and where the seller can be affected (the *Constant 1.0* policy). In the latter case, also the best of *Simple* and even *Adaptive* strategies show improvement. This is what hypothesis 26 predicted, so we can accept it.

Our fourth topic of interest are the effectiveness of future offers, specifically whether they improve the buyer's performance or not. To this end, tables 10.11 and 10.12

Affected Parties	Constant 0.0										Constant 1.0									
	Effect Probability										Effect Probability									
	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
Quality vs. Random Opponent Selection																				
Buyer	+++	+++	+++	+++	+++	+	-	-	-	− − −	+++	+++	+++	+++	+++	+++	+++	+++	+	-
	+++	+++	+++	+++	+++	+++	++	+	++	-	+++	+++	+++	+++	+++	+++	+++	+++	++	-
	+++	+++	+++	+++	+++	+++	+++	+++	+	-	+++	+++	+++	+++	+++	+++	+++	+++	++	-
Seller	+++	+++	+++	+++	++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++
	+++	+++	+++	+++	+++	+++	+++	+++	+++	-	+++	+++	+++	+++	+++	+++	+++	+++	+++	++
Both	+++	+++	+++	+	++	++	++	+++	+	+	+++	+++	+++	+++	++	++	+	+	-	-
	+++	+++	+++	+++	+++	+++	+++	+++	++	+	+++	+++	+++	+++	+++	++	-	-	+	-
	+++	+++	+++	+++	+++	+++	+++	+	-	-	+++	+++	+++	+++	+++	-	-	+	-	-
Expected Utility vs. Quality Opponent Selection																				
Buyer	-	-	-	−	-	-	-	-	+++	+++	+++	-	+	+++	+++	+++	+++	+++	+++	-
	-	-	-	-	-	-	-	-	-	-	-	-	+	+++	+++	+++	+++	+++	+++	++
	-	-	-	-	-	-	-	-	-	-	++	-	-	+	++	+	+	++	++	-
Seller	-	-	-	+	-	++	-	-	-	-	-	-	-	-	+	-	-	-	-	-
	-	-	+	++	+++	+++	+++	+++	+++	+++	-	-	-	-	-	-	-	-	-	-
	-	-	-	-	-	-	-	-	-	-	++	+	-	-	-	-	-	-	-	-
Both	-	-	-	-	-	-	-	-	-	-	+++	-	-	-	-	+	-	-	++	+++
	-	-	-	-	+	+	-	-	+	-	-	-	-	-	-	-	-	-	-	++
	-	-	-	-	-	-	-	-	-	-	++	-	-	-	-	-	-	-	-	-

TABLE 10.10: Effect of Opponent Selection (Adaptive Counter Tactic with Future Offers) (Hypothesis 26). Best of Simple on the top row, Analytic in the middle and Adaptive on the bottom row.

show the results with the *Random Counter* and *Adaptive Counter* tactics, respectively. Basically there is a clear and definite advantage associated to the use of future offers in a clear majority of cases for all the best concurrency strategies. There are three types of exception. First, the very high buyer effect probabilities with the *Constant 1.0* decommitment policy. This is because almost all contracts will be made after the last matching, when future offers are zero anyway. This lack of performance improvement is present in all cases where the *Constant 1.0* fee is used, when the buyer alone can be affected. There are some signs of it also when both can be affected and a sophisticated opponent selection strategy is used.

Second, when the seller is affected and the *Constant 1.0* decommitment policy is used, we get sporadic differences with the *Quality* or *Expected Utility* opponent selection for the *Random Counter* and with the *Expected Utility* opponent selection for the *Adaptive Counter* tactic cases. This is because these opponent selection strategies order the opponents in descending order of expected utility, so the differences between the different opponents are very small, which usually means that only a few negotiations are necessary. Also, the best contracts are often formed very close to the beginning of the run. This means that all strategies take the first contract they can and the expected utilities are very close to each other. The thresholds will not be able to help here because they would only remove opponents that would not bring any result or at least slightly inferior results. The *Adaptive* concurrent strategy will often be able to find the best deal already with the *Quality* opponent selection even with the *Adaptive Counter* negotiation tactic and there is often no improvement there.

Third, a similar effect can also be detected with the same opponent selection strategies where the *Constant 0.0* policy is used and the buyer alone is affected. This is because the first successful negotiation is likely to get a very similar outcome to any later negotiation. This is because with very high buyer effect probabilities many of the contracts are never going to performed and the remaining few are often going to be quite good given the opponent selection strategies. With the less sophisticated opponent selection, the first contracts might not be that good, which is why the future offers are useful. And with lower probabilities of buyer failure, there are more contracts performed so also smaller differences can become significant. Unlike in other cases, the case with the best *Simple* strategy, the *Quality* opponent selection and the *Random Counter* tactic has no improvement in small effects either. This is because the first offer that the *Simple 1* strategy can find even without the future offers is often quite good (very high quality), because it may take a while to find that contract and new and better opponents may arrive before the contract is found so all that the future offers achieve in this setting is to make it possible to get a very similar performance with almost any *Simple* strategy. However, with the *Expected Utility* opponent selection, a suitable contract is often found immediately even without future offers, which means that it can sometimes be merely good (not excellent) and here the future offers may encourage the buyer to wait a bit which means that the future offers may be able to help. All these findings are consistent with hypothesis 27 and we can therefore accept it.

And finally, we will discuss the differences between the two negotiation tactics. The differences between these negotiation tactics with different concurrency strategies in different settings are in table 10.13, where the top row of each setting describes the differences for the best *Simple* strategy, the middle row for the *Analytic* strategy and the bottom row for the *Adaptive* concurrency strategy. Here, the plusses refer to the cases where the strategy using the *Adaptive Counter* tactic outperforms the strategy using the *Random Counter* tactic and minuses denote the contrary situation. The number of plusses again indicate the significance level of the statistical test. We do have plusses, minuses and hyphens, so all possible cases are present. As can be seen, the intensity of the plusses increases with the opponent selection strategies' sophistication. For example, when the seller is affected and the *Constant 0.0* fee is used, there is a difference in most cases only with the *Adaptive* strategy when the *Random* opponent selection is used, with both the *Simple* and *Adaptive* when the *Quality* opponent selction is used and in all three, when the *Expected Utility* opponent selection is used. This pattern also occurs elsewhere.

Affected Parties	Constant 0.0										Constant 1.0									
	Effect Probability										Effect Probability									
	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
Random Opponent Selection																				
Buyer	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	++	-	-	-
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	-	-	-
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	-	-
Seller	+++	+++	+++	+++	++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++
Both	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	++	-
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++
Quality Opponent Selection																				
Buyer	+++	-	-	-	-	-	-	-	-	-	-	+++	+++	+++	+++	+++	+++	+	-	-
	+++	+++	-	+++	+	-	+	+	+	-	+++	+++	+++	+++	+++	+++	+++	+++	+	-
	+++	+++	+++	+++	+++	+++	+	-	-	+	+++	+++	+++	+++	+++	+++	+++	+++	-	++
Seller	++	+++	+++	+++	+++	+++	+++	+++	+++	+++	-	-	-	-	-	-	-	-	-	-
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	++	+	-	+	++	-	-	-	-
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	++	-	-	-	-	++	++	-	-
Both	+++	++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+	++	+++	+	+++	+	++	-
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+	++
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	-	-
Expected Utility Opponent Selection																				
Buyer	+++	+++	++	++	++	+	+	-	-	-	+++	+++	+++	+++	+++	+++	+++	+++	-	-
	+++	+++	+++	++	-	++	-	++	-	-	+++	+++	+++	+++	+++	+++	+++	+++	++	-
	+++	+++	+++	+++	+++	++	+++	++	-	-	+++	+++	+++	+++	+++	+++	+++	+++	-	+
Seller	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	++	-	-	++	-	-	-	-	-
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	++	+	-	-	-	-	-	-	-
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	++	-	+	++	-	-	-	-
Both	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+	-	-
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	-	-
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	++	-

TABLE 10.11: Effect of Future Offers (Random Counter Tactic) (Hypothesis 27). Best of Simple on the top row, Analytic in the middle and Adaptive on the bottom row.

Affected Parties	Constant 0.0										Constant 1.0									
	Effect Probability										Effect Probability									
	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
Random Opponent Selection																				
Buyer	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	-	-	-	-
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	-	-	-
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	-	-
Seller	+++	+++	+++	+++	++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	-	-	-
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+	-	-	-	-
Both	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	++	+++	+++	++	++	+
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	++	++
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	++	++	-
Quality Opponent Selection																				
Buyer	+++	+++	+++	+++	+++	+++	+++	+++	+++	-	+++	+++	+++	+++	+++	+++	+++	-	+	-
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	++	-	-
	+++	+++	+++	+++	+	+++	+	+++	+	+++	+++	+++	+++	+++	+++	+++	+++	+++	-	-
Seller	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	-	++	-	++	-	-	-	-	+
Both	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+	-
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	++
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	++	-
Expected Utility Opponent Selection																				
Buyer	+++	+++	++	++	+	-	-	-	-	-	+++	+++	+++	+++	+++	+++	+++	+++	-	-
	+++	+++	+++	+++	+++	+++	+	-	-	-	+++	+++	+++	+++	+++	+++	+++	+++	-	-
	+++	+++	+++	+++	+++	+	++	-	-	-	+++	+++	+++	+++	+++	+++	+++	+++	+	-
Seller	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	-	-	-	-	-	-	-	-	-
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	++	-	++	-	-	+	+	-
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	++	-	-	-	-	-	++	-
Both	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	-	-	-
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+	+
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	-	-

TABLE 10.12: Effect of Future Offers (Adaptive Counter Tactic) (Hypothesis 27). Best of Simple on the top row, Analytic in the middle and Adaptive on the bottom row.

Affected Parties	Constant 0.0										Constant 1.0									
	Effect Probability										Effect Probability									
	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
Random Opponent Selection																				
Buyer	+++	+++	+	+++	++	-	++	-	-	-	−−−	−	−	−−	−−	−−−	−−	-	-	-
	+++	+++	+++	++	+++	+	++	++	-	-	+++	+	-	-	-	-	-	-	-	-
	+++	+++	+++	+++	++	-	-	+	+++	-	+++	+++	+	+++	+++	+++	+++	+++	+	-
Seller	+++	++	-	-	-	-	-	-	-	-	−−−	−−−	−−−	−−−	−−−	−−−	−−−	−−−	−−−	-
	+++	+++	-	-	-	-	++	-	-	-	+++	+++	+++	++	+	-	++	++	-	-
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	++	++	-	++	-	+
Both	+++	+	+++	+++	-	-	+	-	-	-	−−−	−−−	−	−	-	-	-	-	-	-
	+++	+	+++	+++	-	-	-	-	++	-	+++	++	-	+	-	-	-	-	-	-
	+++	+++	+++	+++	+++	++	+	++	+	+	+++	+	-	-	+	+	-	-	+	-
Quality Opponent Selection																				
Buyer	+++	+++	+++	++	-	-	-	-	−−−	−−−	+++	+++	+++	++	++	++	-	-	-	-
	+++	+++	+++	+++	++	+	-	-	-	-	+++	-	-	-	-	-	-	−	-	-
	+++	+++	+++	+++	+	++	-	+	-	-	+++	-	-	-	+	+++	+++	+++	++	+
Seller	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+	+	-	-	-	-	-	-
	+++	+++	++	-	-	-	-	++	-	-	+++	+++	++	+	-	-	+	-	-	-
	+++	+++	++	+	++	+++	+++	+++	+++	+++	+++	+++	+++	-	+	++	-	-	-	-
Both	+++	+++	+++	+++	+++	+	+++	+++	+	+	+++	+	+	+	-	-	-	-	-	−
	+++	+++	+	-	-	-	-	-	-	-	+++	+	-	-	-	-	-	-	-	-
	+++	+++	-	-	+	+++	-	++	+	-	+++	+	+	-	-	-	-	+	-	-
Expected Utility Opponent Selection																				
Buyer	+++	+++	+++	-	-	-	-	-	-	-	+++	+	-	+	+++	+++	+++	+++	+++	+++
	+++	+++	+++	+++	+++	++	++	-	-	-	+++	+	+	+++	+++	+++	+++	+++	+++	+++
	+++	+++	+++	-	++	+	-	-	-	-	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++
Seller	+++	+++	+++	+++	++	++	++	+++	+++	+++	+++	-	-	-	-	-	-	-	-	-
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+	-	-	+	-
	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	++	-	-	+	+	-
Both	+++	+++	+	-	-	-	-	+	-	+++	+++	+++	-	-	-	-	-	+	+	+++
	+++	+++	+++	++	++	++	++	++	++	-	+++	++	++	+++	+	-	+	+	++	+
	+++	+++	+++	++	+	+	++	+	+	-	+++	+++	+++	++	+	++	++	+	-	-

Table 10.13: Adaptive Counter vs. Random Counter (Hypothesis 28). Best of Simple on the top row, Analytic in the middle and Adaptive on the bottom row.

This is because the strategies are ever more able to access opponents with the *Exponential Time-Dependent* tactic and they can, therefore, exploit them. This is the reason also for why the *Adaptive* concurrency strategy usually is the first one to get the advantage and also to why there is less difference with high buyer effect probabilities.

The only exception to the *Adaptive Counter*'s superiority is in cases where the *Random* opponent selection is used together with the *Constant 1.0* decommitment policy. Here, the *Random Counter* tactic is able to outperform the *Adaptive Counter* in many cases when the best *Simple* concurrency strategy is used. However, there is a relative simple explanation to this. This is because the risk of getting into extra contracts is lower (even with future offers) and the *Simple* strategies with more opponents are practical. Also the effect is especially clear when the seller is affected and, therefore, when a quick contract is needed and the threshold at first matching might not be very high (there are plenty of opponents). Only the *Simple* strategies will be able to benefit because the advanced strategies are already able to avoid extra contracts in both cases. Given that it is good to look for more options when the variation is high, the simple strategies do well here. Also when the buyer is affected, the lower chance of extra contracts helps. All this is consistent with hypothesis 28 and we can accept it.

Affected Parties	Constant 0.0										Constant 1.0									
	Effect Probability										Effect Probability									
	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
Random Opponent Selection																				
Buyer	-	−−−	−	-	-	-	-	-	-	-	-	-	-	-	+	++	-	-	-	−
Seller	-	-	-	-	-	-	-	-	++	-	-	+++	-	-	-	++	-	-	+	+++
Both	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
Quality Opponent Selection																				
Buyer	-	-	-	-	-	-	-	-	-	-	+++	++	++	+	+	+	-	-	-	−
Seller	-	+	-	+++	++	+	++	++	++	+++	+++	+	+++	+++	+++	++	+++	+++	+++	+++
Both	-	-	-	+	-	-	-	-	+	-	+++	++	+	-	-	-	-	-	-	-

TABLE 10.14: The performance of the *Analytic* concurrency strategies under full and no tactic information (the *Random Counter* tactic with future offers) (Hypothesis 30.a).

10.2.3.3 No Tactic Information with Future Offers

When the buyer does not have any idea what tactic an individual seller is employing, its task becomes more complicated. We start with the case where the *Random Counter* tactic is used at the `Negotiator` level. The results with the *Random* and *Quality* opponent selection are in figures 10.16 and 10.17 respectively (performance in the full information case is included for comparison). The detailed analysis of the relevant performances of full and no tactic information are in table 10.14. The plusses in that table refer to the cases where the *Analytic* strategy in possession of full tactic information is able to beat the same strategy with no tactic information and the number of plusses tell the significance level of this difference.[34] We did two-sided tests because we expected that we might also get some cases where the information might actually be counterproductive. We do get two such cases when the buyer is affected with probability of 0.1 or 0.2 and the *Random* opponent selection and the *Constant 0.0* policy are used.

The performances of the *Analytic* and the best *Simple* strategy are compared in table 10.15. As can be seen, there is very little difference between the *Analytic* strategy and the best *Simple* strategy when the *Random* opponent selection is used. There seems to be difference only when the seller effect is very likely and the fee is zero and also when the effect probability is zero and the fee is one. This is because the best *Simple* strategies are over-eager to get into contracts (see no risk in that) but there is a risk, having to decommit from multiple contracts and the buyer using the *Analytic* strategy is able to avoid this. When the effect probability

[34] As in the earlier cases, one plus refers to the $p < 0.05$ level, two plusses to the $p < 0.01$ level and three plusses to the $p < 0.001$ level.

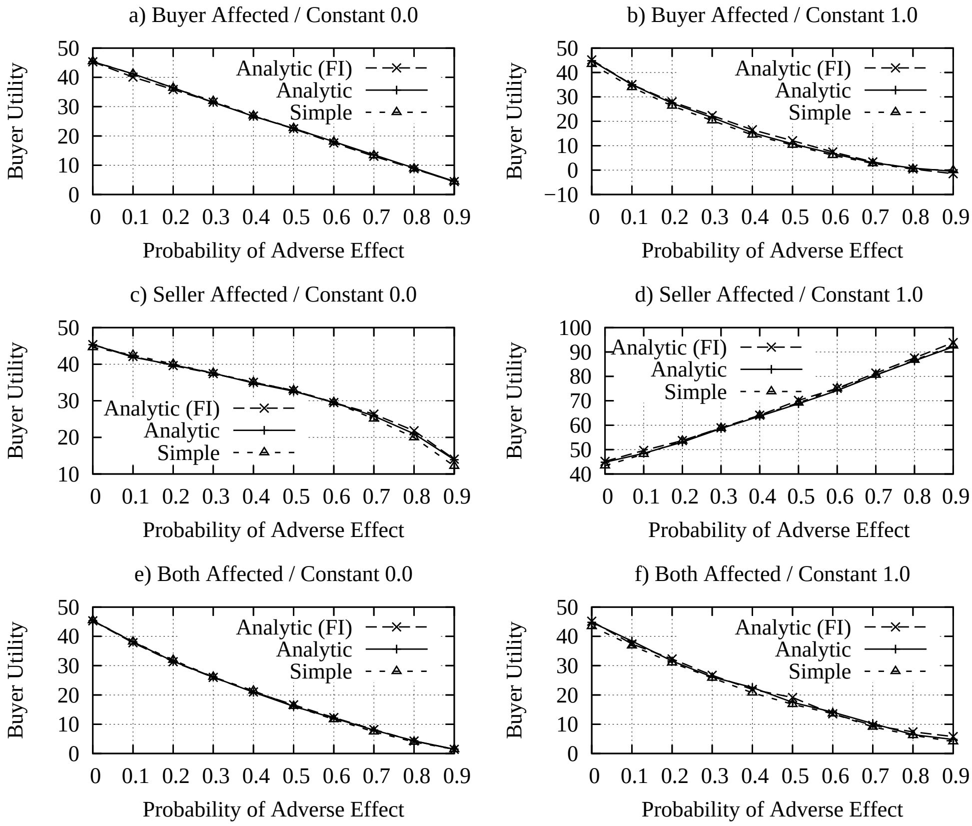

FIGURE 10.16: The performance of different concurrency strategies with the *Random Counter* tactic, the *Random* opponent selection strategy, future offers and no tactic data (Hypothesis 30). The performance of the *Analytic* strategy with full tactic information (FI) is also given for comparison.

increases, also the *Simple* strategies become more careful and the *Analytic* strategy is unable to beat the best *Simple* strategies.

When the *Quality* opponent selection is used, things change. Now, all opponents are high quality and around half of them will also accept an offer made by the *Random Counter* tactic. The buyer just does not know with which opponents the negotiation will succeed and with which opponents it will fail. When there is a significant benefit in waiting, the *Simple* strategies will either waste a lot of perfectly good opponents (the *Simple* strategies with a high number of negotiations) or are unable to find good opponents when the time comes (the *Simple* strategies with a low number of negotiations). The *Analytic* strategy will be able to wait until there are probably good opponents (those it can succeed with) among the sellers and it is able to estimate a reasonable number of opponents to negotiate with to get a result with a reasonable risk of getting into extra contracts. So, when the *Constant 0.0* decommitment policy is used, this means that when the seller can

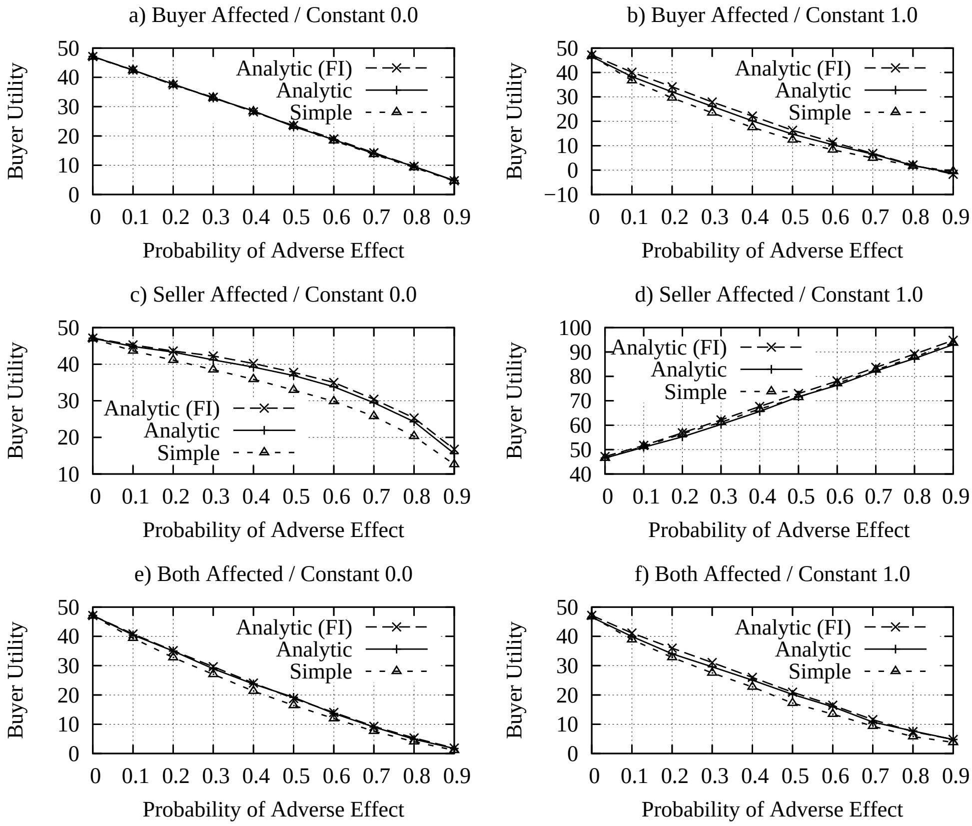

FIGURE 10.17: The performance of different concurrency strategies with the *Random Counter* tactic, the *Quality* opponent selection strategy and future offers but with no tactic information (Hypothesis 30). The performance of the *Analytic* strategy with full tactic information (FI) is also given for comparison.

be affected (waiting is useful), the *Analytic* strategy will be able to beat all *Simple* strategies. If there are no adverse effects, good contracts can be found at any time and extra contracts can be decommitted from for free, the best *Simple* strategies will be able to do as well as the *Analytic* strategy.

Now, when the *Constant 1.0* decommitment policy is used, the situation is more complex. The *Analytic* strategy is able to beat the competition in most cases where the waiting is useful (when the buyer is affected). When only the seller is affected, however, the best strategy is to get a contract as soon as possible and, here, also the *Simple* strategies can be successful. In fact, the *Simple* strategies seem to be able to beat the *Analytic* strategy in some cases. It seems that with no information about the opponents' tactics, the *Analytic* strategy somewhat underestimates the risk of getting into multiple contracts. It often starts with three negotiations in the first matching and it has to decommit (on average) almost two contracts in every hundred runs , whereas the best *Simple* strategy is *Simple 1* which does not need

Affected Parties	Constant 0.0										Constant 1.0									
	Effect Probability										Effect Probability									
	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
Random Opponent Selection																				
Buyer	-	-	-	-	-	-	-	-	-	-	+++	-	+	+	-	-	-	-	-	-
Seller	-	−	-	-	-	-	-	+	++	+++	+++	-	-	-	-	-	-	-	-	-
Both	-	-	-	-	-	-	-	+	+	-	+++	+	-	-	-	-	-	-	-	-
Quality Opponent Selection																				
Buyer	-	-	-	-	-	-	-	-	-	-	−−	+++	+++	+++	+++	+++	+++	+++	-	−
Seller	-	+++	+++	+++	+++	+++	+++	+++	+++	+++	−−	-	-	−−	-	−	-	-	-	−
Both	-	+++	+++	+++	+++	+++	+++	+++	+++	+++	−−	-	+	+++	++	+++	+++	+	++	-

Table 10.15: The performances of the advanced concurrency strategies as compared to the best *Simple* concurrency strategy (the *Random Counter* tactic with no tactic information but with future offers) (Hypothesis 30.b). The *Analytic* vs. the Best of the *Simple* on the top row and the *Adaptive* vs. the Best of the *Simple* on the bottom row in each case.

to decommit from extra contracts at all. This comes from the fact that the two tactics that are the *Random Counter* can succeed with have very different contract time patterns. The negotiations with the *Random* tactic can succeed at any time, but earlier times are slightly more probable, whereas with the *Exponential Time-Dependent* tactic, the negotiation can often take a while. We also assumed for the sake of simplicity (see section 10.1.2.2) that the *Exponential Time-Dependent* tactic always succeeds at the last turn, which is not necessarily the case with tactics that concede a lot quickly and also with the slightly less quick conceders, the buyer's offer may be acceptable sooner than in the last turn (because it is higher than the reservation price). When the seller tactic is known, such small differences do not play a significant role, but when uncertainty increases, this effect is enough to make the *Analytic* strategy lose to the *Simple 1* strategy. All these findings are consistent with hypothesis 30 and, therefore, we can accept it.

The cases where the *Adaptive Counter* tactic is used at the `Negotiator` level are in figures 10.18 and 10.19. It is very clear from these figures that the lack of tactic information decreases the performance of the buyer strategies significantly in almost all cases. An interesting observation is a sudden increase of performance when the buyer is affected with the probability 0.1 and the *Constant 1.0* decommitment policy is used. This effect occurs with both the *Random* and *Quality* opponent selection and is because the large decommitment fee makes the *Adaptive Counter* tactic use lower offers. This, in turn, is because making offers that are always going to be accepted does not give the buyer a very good expected utility when there is a risk of expensive decommitment. A lower offer that gives the buyer higher expected utility in case of success will be more interesting because of the higher expected utility itself and because it decreases the chance of having

Affected	Constant 0.0										Constant 1.0									
Parties	Effect Probability										Effect Probability									
	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
Random Opponent Selection																				
Buyer	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+	++	-	-	-	−−	-
Seller	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++
Both	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++
Quality Opponent Selection																				
Buyer	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	-	−	−−	-
Seller	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++
Both	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	++

TABLE 10.16: The performance of the *Analytic* concurrency strategies under full and no tactic information (the *Adaptive Counter* tactic with future offers) (Hypothesis 29.a).

to pay the high fee (there is a positive probability that no contract will be formed at all giving zero profit). All this means that the *Adaptive Counter* will be more risk-seeking when only the buyer will be affected. When the seller is affected there is no similar incentive because there is very little risk of negative utility for the buyer in case of a contract. This also holds when both parties can be affected.

Now, in table 10.16, we have a detailed analysis of relative performance of the *Analytic* concurrency strategy both under full and no information. As can be seen from this table, the difference is very clear in almost all cases. The only exceptions are associated with the *Constant 1.0* policy when the buyer is affected. Here, with high effect probabilities, with both the *Random* and *Quality* opponent selection, as we just discussed, the *Adaptive Counter* tactic will take more risks and this means, it will do very well, even without tactic information. There are even a couple of settings where the *Analytic* strategy does better without than with the tactic information. As explained earlier (section 10.1.3), such an effect is possible because the thresholds demand higher quality without than with the tactic information and because this may mean the *Analytic* strategy negotiates more. Of course with very high buyer effect probabilities this means the negotiations take place very near the deadline and it simply means that unlike with full information, the *Analytic Counter* without the tactic information will demand a higher expected utility which means that it will fail negotiations with behavioural tactics, an option it does not have with full information. All these observations are consistent with hypothesis 29.a.

Now, the other question we will discuss in relation to the *Adaptive Counter* tactic is how the *Analytic* strategy operates compared to the *Simple* strategies. The summary of this performance is in table 10.17. We start from cases with the *Constant 0.0* policy. With the *Random* opponent selection, there are some cases where

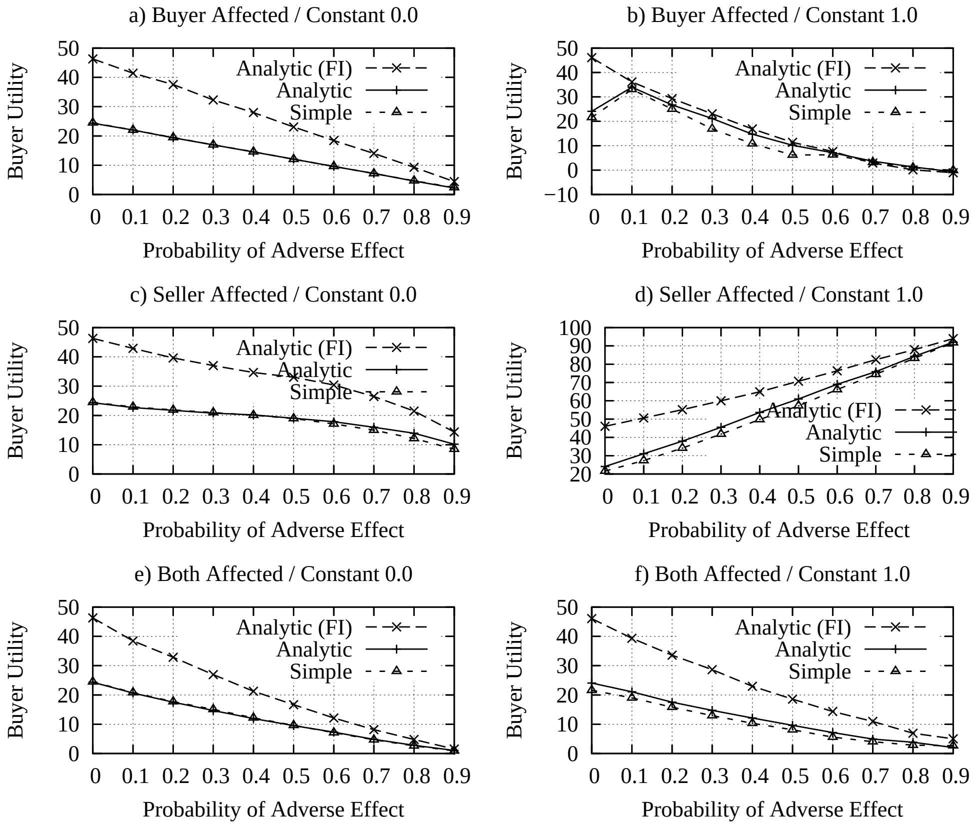

FIGURE 10.18: The performance of different concurrency strategies with the *Adaptive Counter* tactic, the *Random* opponent selection strategy, future offers and no tactic data (Hypothesis 29). The performance of the *Analytic* strategy with full tactic information (FI) is also given for comparison.

the best *Simple* strategies beat the *Analytic* strategy with low effect probabilities. As before, this is because the *Analytic* strategy does not negotiate except when it knows it will be successful, which means that it will encounter some opponents repeatedly and therefore it will meet fewer new opponents than some of its *Simple* competitors. However, when the waiting starts to pay off (when the seller effect probability increases), the *Analytic* strategy will be able to beat all the *Simple* strategies. With the *Quality* opponent selection, the number of opponents loses its significance and all the best opponents are encountered by all strategies. This means that the *Analytic* approach to wait until the time is right to negotiate wins the day when the seller or both parties can be affected. As in the earlier settings, there is no such advantage when the buyer is affected because delaying the contract does not have an effect on the buyer's effect probability and there is no difference between having and not having a contract when and if the decommitment occurs. All this is very clear. In large part this is because the *Analytic Counter* tactic will use offers that will be accepted by any seller and therefore only the quality of the

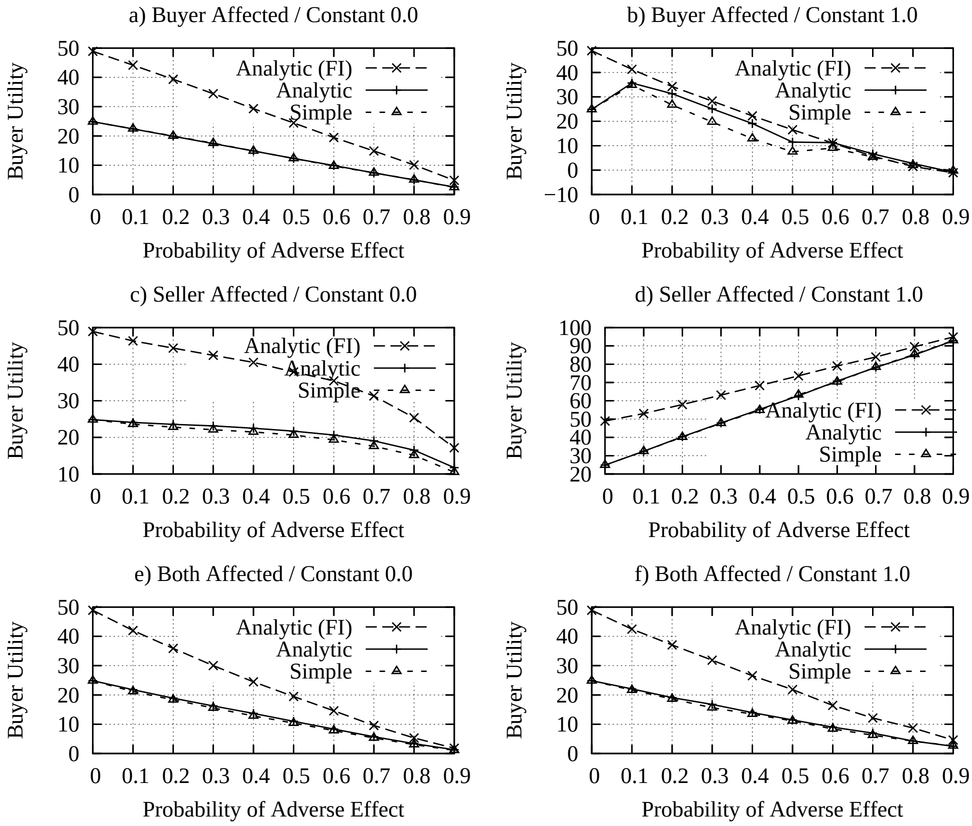

FIGURE 10.19: The performance of different concurrency strategies with the *Adaptive Counter* tactic, the *Quality* opponent selection strategy and future offers (Hypothesis 29). The performance of the *Analytic* strategy with full tactic information (FI) is also given for comparison.

opponent and the time of contract (when the seller is affected) can influence the results.

The situation is more complicated when the *Constant 1.0* fee is used because here, some negotiations also fail. However, there are no real surprises here. The *Analytic* strategy is able to beat its *Simple* competitors in most cases with the *Random* opponent selection. The only exceptions are the very high effect probabilities. When seller or both parties are affected, the *Adaptive Counter* makes offers that are always accepted. Because of the high decommitment fee, the *Analytic* concurrency strategy usually negotiates with only the first opponent that the future offers do not remove and the best *Simple* strategies do the same and the best *Simple* strategies often do something very similar. They are a bit less accurate with this[35] but with very high effect probabilities only a handful of contracts are

[35]This is because in the n opponents, the *Simple* n strategy encounters, there might be also zero or more than one opponent that exceeds the threshold and, therefore, the *Simple* n strategy may be left with no or too many contracts.

Affected Parties	Constant 0.0										Constant 1.0									
	Effect Probability										Effect Probability									
	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
Random Opponent Selection																				
Buyer	– – –	-	-	-	-	-	-	-	-	-	+++	-	++	+++	+++	+++	-	-	-	–
Seller	– – –	––	-	-	-	-	++	-	+++	+++	+++	+++	+++	+++	+++	+++	++	++	+	-
Both	– – –	-	-	–	-	-	-	-	-	++	+++	+++	++	+++	++	+	+	-	-	-
Quality Opponent Selection																				
Buyer	+++	-	-	-	-	-	-	-	-	-	+++	-	+++	+++	+++	+++	+++	++	-	–
Seller	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	+++	-	-	-	-	-	-	-	-	-
Both	+++	+++	+++	+++	+++	+++	++	++	++	+	+++	-	-	++	-	-	-	-	-	-

Table 10.17: The performances of the advanced concurrency strategies as compared to the best *Simple* concurrency strategy (the *Adaptive Counter* tactic with no tactic information but with future offers) (Hypothesis 29.b). The *Analytic* vs. the Best of the *Simple* on the top row and the *Adaptive* vs. the Best of the *Simple* on the bottom row in each case.

ever performed and this difference drowns in the noise. When only the buyer is affected, the buyer is taking chances in its negotiations (demanding a lower price) and it fails quite often with the negotiations it starts. But, as in the other cases, there are less contracts and less possibilities for making a difference. With the extreme buyer effect (0.9), we also (as in the cases with full tactic information) get the slight overeagerness to negotiate (due to inaccuracies in expected utility estimates as discussed in earlier settings) and the flexibility of the *Analytic* means it is better able to find these problem cases. This explains why the best *Simple* strategies are able to beat the *Analytic* strategy there.

When the *Quality* opponent selection is used, there is a clear and sustained advantage for the *Analytic* strategy only when the buyer is affected. This is very simply because in other cases, the opponents make offers that are always accepted and therefore having many negotiations is not that good an idea. All strategies will therefore negotiate with the first opponent and make it an offer it eventually accepts. There is no room for improvements through using the *Analytic* strategy. All this is consistent with hypothesis 29.b and we can therefore accept the whole hypothesis 29.

10.3 Summary

In this section, we discussed the `Controller` level and described in detail our model for this level. We considered for the first time, opponent selection and concurrency control (choosing the number of negotiations), introduced several strategies for both and defined the clear rules for interaction between these two and other strategies (contribution **C7**). We also successfully used empirical data on the previous runs to estimate the offers the buyer would be encountering later (contribution **C9**). We tested several hypotheses about these strategies empirically and discovered that the opponent selection, the concurrency strategies and considering future offers can have a significant impact on the performance of the buyer agent. Generally, the more sophisticated strategies did better than their simpler counterparts, but this was not always given and there were also cases where the simpler strategies outperformed the sophisticated ones.

We also discovered that operating in the presence of incomplete information has a strong impact on the performance of the buyer agent when the *Analytic Counter* tactic is used at the `Negotiator` level and much less effect when the *Random Counter* tactic is used (contribution **C7**). Although the *Analytic Counter* tactic would certainly benefit from the *Adaptive* concurrency strategy, which would have allowed it to take into account the other negotiations, the *Adaptive Counter* would still require quite detailed information about the opponent population to function effectively. The simple counter tactics, such as *Random Counter*, do need significantly less information and are less affected by the lack of information. Although as we discovered, the *Simple* strategies might in some cases still be advisable.

Chapter 11

Coordinating Concurrent Negotiations: The `Coordinator` Level

The top level of our model consists of the `Coordinator` that coordinates between negotiations on different services. The aim here is to discuss our contributions **C8** (managing basic interrelations between negotiations). We do this by discussing substitutes and complements and by developing some basic strategies for managing them in a situation where the buyer may be adversely affected and change its mind about needing the service during the negotiation. By so doing, we wish to show that the `Coordinator` level is an essential part of a concurrent bilateral negotiation agent, if a buyer agent is simultaneously engaged in many interrelated negotiations and not just singular acquisitions. As discussed in the literature review (section 2.3.4), the existing work on concurrent bilateral negotiation has been lacking in this respect. Naturally, future work can extend these fundamental strategies into more complex and efficient ones, but the aim here is to lay the foundation in this area.

We start by discussing how the `Coordinator` is constructed and how it interacts with its environment (section 11.1). We then discuss our experiments and results (section 11.2), and we conclude the chapter with a summary of our findings (section 11.3).

11.1 Architecture of the `Coordinator`

The `Coordinator`'s task is to coordinate negotiations on different services. It will start new and end old `Controller`s as necessary. The `Coordinator` gets its targets from the planner or human user, both of which are out of scope for this work. Here, it is simply assumed that the `Coordinator` just gets the necessary information (the markets and substitute/complements). In a complete system, the `Coordinator` might give them intermediate results and obviously the planner or human user could at any time change the requirements, but we will not consider such changes in this work. Instead, we assume the `Coordinator` simply gets its targets and autonomously tries to fulfil those targets to the best of its abilities.

In the work we discuss here, we have between two and nine separate service markets. The reason for the minimum (two) should be self-evident, since it is the smallest number of markets above one. The maximum (nine) was selected because we wanted to have more than a couple of settings and because eight seemed quite reasonable an amount for the graphs. In these eight settings, the `Coordinator` will try to achieve one of two goals depending on the setting:

- In the *Complements* setting, the goal is to get one of each and every service or none of the services at all. An example of complementary services would be a trip from A to B, say from Southampton, UK, to Tokyo, Japan. The trip consists of a trip from Southampton to Heathrow Airport, a flight from Heathrow to Narita Airport near Tokyo and a trip from Narita to Tokyo. All three partial services are needed in order to get from Southampton to Tokyo, so the services are complementary. A trip from Narita to Tokyo is not that useful if we cannot get to Narita and so on.

- In the *Substitutes* setting, only one of the services is required (and it does not matter which one). Using the same example, there are substitute services for the trip from Southampton to Heathrow. Since here we assume that the options are pretty much identical in terms of expected price and utility, we might not consider taking a taxi but instead we could consider taking a coach directly from Southampton to Heathrow or taking a train to Woking and a bus from there to Heathrow. In terms of price and travel time these options might be roughly equal. On the other hand, we need to make the trip only once, so if the agent is negotiating with both the coach and train company at the same time, it needs to be careful not to book both but only one of the services.

In our experiments, the target is always one or the other (getting one of all services or getting only one service). In other situations, the interdependencies between different negotiations might of course be more complicated. This could mean that there are combinations of complements and substitutes in the same problem. An earlier example can be used here again by combining the two cases. The agent could need to find a way to get a person from Southampton to Tokyo and each leg could be done in many different ways. Also there might be other airports the agent might consider (either in Japan or in UK) and this would mean completely new sets of services for the other legs of the journey. So, for example, it might be an option to fly to Singapore and change planes there. Here, we would again get complements from getting to the airport in UK, to Singapore, from there to Japan and from airport to Tokyo and this would be a subsitute to a more direct trip via Heathrow. In cases such as these we might even have a hierarchy of `Coordinator`s, where the lowest level `Coordinator`s, for example, would try to find the best deals for different possible legs of the journey and the higher levels would combine these to the whole journeys and the top level would then choose which of the whole journeys to take. However, as already mentioned, this type of problem is beyond the scope of our work.

Moreover, for the sake of clarity, we have made some additional simplifications. We assume that matchings in different markets always occur at the same time. This could be, for example, because the turns we use would have a counterpart in reality, for example a hundred turns could be a minute of real-time and matchings could occur once a minute. Or it could be because all the markets in the marketplace are synchronised to have matchings at the same time. In either case, we just assume the `Coordinator` can make its decisions involving the number of negotiations in each market at the same time. This is to make the coordination problem somewhat easier, without losing anything essential to the problem itself.

We also assume that all services in all markets are identical in terms of the expected number of entries, quality distribution, utility functions, the expected utilities, and so on. In the case of substitutes, this means that the different service markets provide *perfect* substitutes meaning that at the outset, the buyers are indifferent between the service types. They are equally happy to take a service from any service market (no preferences between services). Between the service markets, the service is selected from the market where the expected (quality) utility is the highest and no market has an advantage in advance. Of course in a more realistic setting, the services are unlikely to be perfect substitutes. Some people prefer riding by coach to taking a train, a train option might involve changes whereas the coach option might not, a train station might be closer to home than the closest coach stop, the actual prices might vary, taking a taxi would also be a possibility (albeit a more expensive one) and so on. So different services would probably have different utility functions and a more realistic system would have to take this into account. However, we do not think that this sort of additional complexity would bring anything essentially new to the problem but instead, in a more simple and standardised setting, we can show the effect the different coordination strategies have on the outcome.

Similarly, we have assumed that the agent needs to decommit from all contracts on services that it will not use and we will only use one decommitment policy, that of *Constant 1.00*. Our choice of decommitment policy was heavily influenced by the fact that an overcompensating decommitment policy will make the successful coordination a more critical task than it would otherwise be. And of course, since the providers of these services are (or can be) separate entities, the agent will have to pay a fee for each and every contract it needs to decommit from. At worst, this can mean decommitting eight contracts because the critical 9th contract failed to materialise. With an overcompensating decommitment policy, this would be a dire situation indeed. Of course, here the decommitment fee is always more than the service price (since the price is always less than 1)[1] and, in practise, the buyer would be tempted to just stay in the contract and pay for the service even if it never plans to use the results for anything. However, we felt that to demonstrate the importance of the coordination, we want to make the payment for failure reasonably large. The differences would be there also in cases where the buyer would just stay in the contract or could get out of the contracts by paying a smaller fee (a zero fee would, however, mean that there was very little need to coordinate), but the differences would be smaller.

Now, after this problem introduction, we will discuss the architecture of the `Coordinator`. First, we will give an overview of the structure of the `Coordinator` and how it interacts with its environment (section 11.1.1). We then discuss specific coordination strategies we are going to use (section 11.1.2) and finally we will discuss how we consider future offers in this environment (section 11.1.3).

[1]The lower levels of the agent ensure that any contract means a positive utility for the buyer, so even if the quality was at maximum (= 1) the price is going to be less than one, so that the buyer is able to get a positive utility out of contract.

11.1.1 Overview

The `Coordinator` level sits on the top of the `Controller` level and all `Controller`s report to the `Coordinator` after they have analysed their respective negotiation situations and have a recommended plan of action in their set of negotiations. The `Coordinator` takes these plans and takes into account the situation in different markets and how these plans go together. It can then make adjustments to the plans. The final plans are sent to the `Controller`s to implement. Each `Controller` will tell its `Negotiator`s what action they should take. Unlike the `Controller` level, the `Coordinator` level automatically starts a `Controller` at each market because this is the only way it can get any information from those markets.

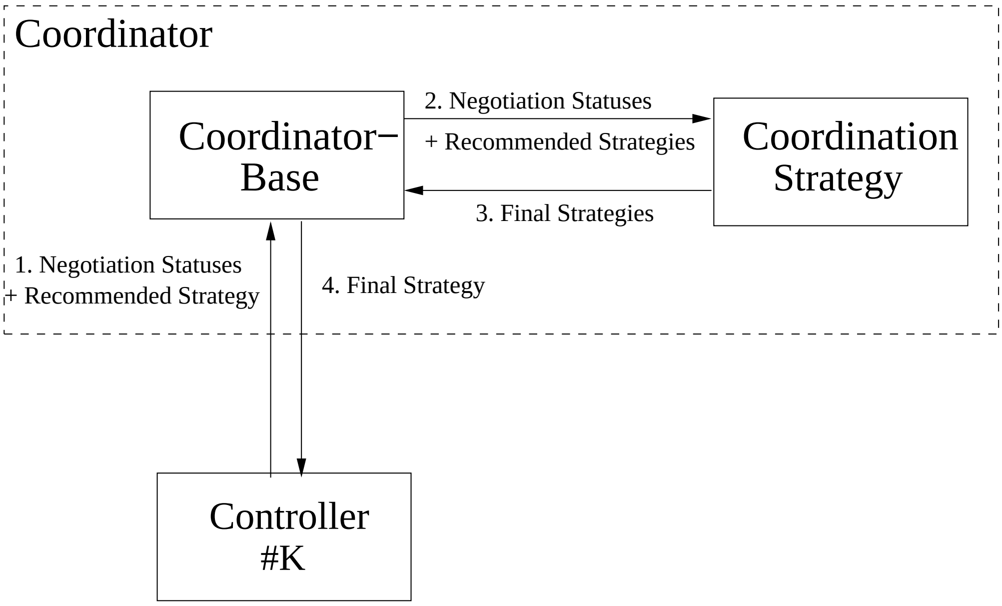

FIGURE 11.1: The `Coordinator`.

In more detail, the structure of the `Coordinator` level itself is very similar to that of the lower levels, especially the `Negotiator` level. Basically, there is a `<CoordinatorBase>` module that handles all the interaction with other levels (here `Controllers`) and an interchangable `<CoordinationStrategy>` module which defines how the negotiations in different service markets are to be coordinated. The `CoordinatorBase` gives the `CoordinationStrategy` the `Controller`'s suggested plans and the `Coordination- Strategy` will give back the adjusted final plans. The `Coordinator` will play a particularly significant role in the opponent selection phase, when the number of opponents are chosen in each market, since this is where the coordination is mostly needed.

Naturally, the `Controller`s keep the `Coordinator` updated on developments in their negotiations, so after a sufficient number of contracts have been secured, the `Coordinator` level will immediately inform all `Controller`s that they should stop their negotiations. This is somewhat trivial and it will happen exactly the same way no matter what coordination policy is used. The differences between coordination policies are in the opponent selection phase alone. To this end, we will now discuss the coordination strategies we have used in our experiments.

11.1.2 Coordination Strategies

These coordination strategies differ in the role they play during the *Opponent Selection* phase, that is when the buyer is choosing who to negotiate with in each market. In contrast, during the *Negotiation* phase (when negotiations are in progress in the markets), the functionality of all coordination strategies is the same. They receive progress reports from the `Controller`s every turn and step in if a sufficient number of contracts (or more) have been formed. In such circumstances, they will end all remaining negotiations in all markets and decommit from extra contracts that are not needed. Otherwise they just approve whatever the lower levels are doing.

As we explained in the beginning of this chapter, we use straightforward coordination strategies. Specifically, we have three such strategies, a default one involving no coordination (the *No* coordination strategy, section 11.1.2.1), one that is geared towards the *Substitutes* case (the *Risk-Averse* coordination strategy, section 11.1.2.2) and one for the *Complements* case (the *All or Nothing* coordination strategy, section 11.1.2.3). We will now discuss each in turn.

11.1.2.1 The No Coordination Strategy

The *No* coordination strategy takes the plans the `Controller`s suggest and accepts them as they are, making *no* changes. Since the `Controller`s are completely unaware of each other, this means that there is no coordination between them. This means that each `Controller` acts independently.

11.1.2.2 The Risk-Averse Coordination Strategy

The *Risk-Averse* coordination strategy is meant for the *Substitutes* setting and it is designed to avoid extra contracts. It does this by negotiating only in one market at a time. At each matching, the *Risk-Averse* strategy will take the recommended strategies from all the `Controller`s and tell the `Controller` that forecasts the highest expected utility to go ahead as planned, while telling all the others to have no negotiations.

This coordination strategy obviously makes it impossible to get into contracts in different markets simultaneously. However, a possible drawback of this strategy is that it may not work very well in environments where there are limited contract opportunities or where the agent needs to find more than one contract. This makes it unsuitable for the *Complements* setting: if the buyer needs a success in 10 markets, negotiating in them one at a time can be very risky. If the probability of the buyer's adverse effect is very high or if the `Controller`s consider future offers (tend to make their move very late in the game), using this strategy would mean automatic failure in cases where many contracts in different markets are required. We therefore use it only in the *Substitutes* setting.

11.1.2.3 The All or Nothing Coordination Strategy

The *All or Nothing* coordination strategy is designed to work in the *Complements* case in environments where there is likely to be times where a contract in all markets is possible around the same time. Basically, at each matching, the *All or Nothing* coordination strategy will see if *all* of its `Controller`s give a very high probability of success. If one or more say that there is a significant chance of failure, the strategy will tell everybody else not to negotiate at all. Only if all are confident of their success, will it give a permission to everybody to go ahead. This, as the *Risk-Averse* strategy before it, is quite conservative a strategy. It will move only if success is all but certain. To enhance this stance even further, it also uses a conservative formulation of its problem. Basically it will calculate the expected utility for going ahead (negotiating in all markets) as follows:

$$EU = P \sum_{i=0}^{n} EU_i - (1 - P)(n - 1)f,$$

where n is the number of markets, P is the probability that a contract will be found in all n markets (this is $P = \prod_{i=0}^{n} P_i$, where P_i is the success probability in market i), EU_i is the expected utility in market i (in case of success) and f is the decommitment fee (in our experiments $f = 1$).

So, the first term calculates the expected utility in case all negotiations succeed and the latter in all the other cases. For simplicity, we have made a pessimistic assumption that all the other negotiations will be successful and therefore the buyer agent will need to decommit from all but one $(n - 1)$ contracts. With the *Constant 1.00* decommitment policy, this formulation means that the success probability will have to be very high indeed ((almost) 1) before the *All or Nothing* coordination strategy gives a go-ahead for its `Controller`s.

Also it should be noted that a very high success probability has to be achieved in all markets at the *same* time (same matching). This can sometimes lead to problems, especially if the markets are very different from each other in terms of timing or if the number of markets is high (the probability that the success probability will be high in all of them at the same time might not be as high as one might wish). It should be remembered that in this simple setting, the `Coordinator` cannot express wishes of success probability to the lower levels so they are stuck with whatever the `Controller`s come up with. It is possible that sometimes a slightly higher price in one market might be acceptable to secure a higher success probability in that market and therefore be successful in getting all contracts. However, that would require a more sophisticated coordination strategy than *All or Nothing*.

11.1.3 Considering Offers in the Future Negotiations

In the work we discuss here, we use future offers exactly like in our experiments in the `Controller` level (see section 10.1.3). That is, each `Controller` considers future offers as if there are no other markets. In some situations, especially in the *Substitutes* setting, the buyer agent would probably benefit from future offers taking into account the greatly expanded opportunities of multiple markets. A deal that is excellent when 250 opponents can be met might be merely good when the number of opponents to choose from increases to 2500, for example. It would probably also decrease the negotiations and therefore the number of decommitments. However, here our task is simply to demonstrate that coordination between different market is important and this can be done also where future offers consider each market in isolation. Also extending the future offers to consider all available markets would be one type of coordination (albeit imperfect) and we wanted to compare the situation with and without coordination. Moreover, such an approach would not be helpful in the *Complements* setting because there the buyer must succeed in all markets (separately) not just one (the market by market approach would be the only possible one).

Given all this, we have therefore *not* expanded future offers to consider all available markets. We will discuss the implications of this decision when we discuss our results.

11.2 Empirical Evaluation

Having explained what the `Coordinator` level looks like, how it works and described our coordination strategies, it is time to perform the empirical evaluation. As before, we start this by discussing our hypotheses (section 11.2.1). We then explain how the experiments were conducted (section 11.2.2) and discuss our results (section 11.2.3).

11.2.1 Hypotheses

We use the best strategies and tactics from the lower level. Thus,we use the *Adaptive* concurrency strategy, the *Expected Utility* opponent selection, the *Adaptive Counter* negotiation tactic (since we are using full tactic information) and the future offers restricted to single markets. The side-effect of our settings and these strategies is that the probability of success is reasonably high in any and all markets. Given the setting, the *Constant 1.0* decommitment policy and the discussion in the previous section, our hypotheses are relatively straight-forward.

We start with the *Substitutes* setting. Here, the *No* coordination strategy runs a high risk of getting into more than enough contracts and this will adversely impact on the buyer's performance. The effect is likely to worsen when the number of markets increases (probability of extra contracts increases). On the other hand, increasing the number of markets means access to more and more opponents, so the *Risk-Averse* negotiation strategy is likely to (slightly) increase its performance as the number of markets increases. The best strategies mentioned above mean that the result in each market is going to be quite good, but if one is in a position to choose the best result from several markets, the results are often better than with one market alone. This effect is of course quite weak because the quite good in one market is going to be very good already and there is limited space for improvement. The effect is also going to be stronger in cases where the buyer's failure probability (chance of adverse effect) is lower because obviously in case of adverse effect even an exceptionally good agreement will not matter. Therefore we contend:

> **Hypothesis 31.** In the *Substitutes* setting, a buyer agent using the *Risk-Averse* coordination strategy will outperform one using the *No* coordination strategy.

Now, the *Complements* setting is likely to be more complicated and simple coordination may even be counter-productive, especially when the number of markets (required contracts) is high. This is because, as explained, the *All or Nothing* coordination strategy will only negotiate if it can be relatively certain that it can get a contract in *all* negotiations in the *same* matching. However, as explained when expected future offers (average later offers) are used as a utility threshold, it means that sometimes starting the negotiations is postponed because opponents with sufficiently high expected utility do not appear. When the number of markets increases, the probability that this happens at least in one market increases. This means that the *All or Nothing* coordination strategy starts negotiations later and some of the best opponents in some other markets may have exited by then, whereas the *No* coordination strategy may negotiate with these opponents whenever they are available. However, when the effect probability increases also the *No* coordination strategy will start negotiations later and later, because the highest utility is often available near the deadline when the probability of adverse effect is at its lowest. Therefore the difference in timing between the two coordination strategies decreases when the effect probability increases.

Another observation is that in a single negotiation, surprising things can happen, for example, the opponent may have a very early deadline. This can mean that a negotiation can fail even if in advance, it seemed quite safe a bet. If you have multiple chances to negotiate, this is probably not going to an issue but when the probability of adverse effect increases, all the best contracts are found near the delivery time, which limits the number of useful negotiations the buyer has and, therefore, increases the chance that the buyer is left without a contract in the end. And again, the more markets you have, the more likely this sort of surprising failure at least in one market is. This means that with very high effect probabilities, the *All or Nothing* strategy may be able to beat the *No* coordination strategy, that will not consider such risks. Thus, we contend:

> **Hypothesis 32.** In the *Complements* setting, with the *Adaptive* concurrency strategy, the *Expected Utility* opponent selection and the *Adaptive Counter* negotiation tactic and no market failures, the *All or Nothing* coordination strategy will be able to beat the *No* coordination strategy only in cases where the adverse effect probability is very high. The *All or Nothing* strategy may even be beaten by the *No* coordination strategy in some cases with very low adverse effect probability.

Now, this would seem to indicate that coordination is a bad idea in some cases. This anomaly occurs because in the setting we have here getting into contracts is usually easy. In other words, the `Controller`s are usually able to ensure success if left to their own devices. However, if they are restricted in the way the *All or Nothing* coordination restricts them, the results are no longer guaranteed or at least not necessarily as good. However, there are situations where a positive result (a contract) cannot be obtained in a market and sometimes this has very little to do with the buyer's actions. To cover this, we introduce a possibility of a *market failure*. This means that with probability z we allow no entries to one of the markets (selected at random), which means that no contracts can be formed in that market. Because in the *Complements* setting, the buyer needs to find a contract in all markets, this means that sometimes (with the probability z) the buyer will be unable to succeed.

The *No* coordination strategy will be defenceless against such a situation. It will enter into contracts in all or at least most of the other markets and will have to decommit from all of these extra contracts later. This will seriously dent its performance. In contrast, the *All or Nothing* coordination strategy will be able to avoid getting into contracts when there is no chance of success, so it will not even negotiate in the market failure cases. It will, however, be as effective as before in those cases where the market failure does not happen. Its performance will, therefore, drop less. This effect should be quite clear already with relatively small failure probabilities. We contend:

> **Hypothesis 33.** In the *Complements* setting, the *Adaptive* concurrency strategy, the *Expected Utility* opponent selection and the *Adaptive Counter* negotiation tactic, the *All or Nothing* coordination strategy will be able to outperform the *No* strategy in all settings with a small amounts of market failures.

11.2.2 Experimental Setup

The experiments were conducted by running the market 100 times and adding the results together. The results here are as in the earlier settings as far as the *Substitutes* setting is concerned. Any extra contracts (beyond 1) are decommitted from and the result is the buyer's utility, so the contract's value minus its price minus fees for any decommitments. In the *Complements* case, however, things are slightly more complicated. Here, if the buyer has a contract from all markets then its utility is equal to the sum of utilities (value−price) for all these contracts minus any decommitments made. This number will therefore increase with the number of markets. With one market, the maximum utility is 0.50, with two, its 1.00 and so on. Since in our experiments we run the market 100 times and add the utilities together, the results discussed in the next section are hundred times as big as the single result (so up to 50, 100, and so on). And if the buyer fails to get a contract from all markets in the end, it will have to decommit from all contract it does have and therefore the utility will be equal to $-nf$, where n is the number of decommitments and f is the decommitment fee. We then repeat this procedure 100 times to get an average result and the variation for 100 repetitions. We will again use two-sided t-tests to see if there are differences to one or the other direction to investigate any differences between the strategies.

We will investigate only the cases where the buyer and the buyer alone can be affected. This is because the other settings would not bring anything essentially new to the problem or at least nothing new to the *coordination* problem, which is in the *Substitutes* case about avoiding extra contracts and in the *Complements* case about securing all or no contracts at all. The case where the seller is affected for example would only bring in cases where the contracts are entered mostly in the first round of negotiations (to maximise the possibility of seller failure). Although this certainly would make coordination very important in the *Substitutes* case (because most contracts happen around the same time), it would not bring anything essentially new to the problem and the same goes for *Complements* case, where we would still have mostly successful contracts. And of course here, it might not even be optimal to avoid getting into contracts if there is a good chance of the seller decommitting and having to pay a huge decommitment fee.

11.2.3 Results

In figures 11.2.a-h, we can see how the *Risk-Averse* coordination strategy performs against *Slave* coordination with varying number of markets. The superiority of the *Risk-Averse* coordination is clear (and statistically significant at the $p < 0.0001$ level) in all cases and the difference between the two strategies increases as the number of market increases. This is mostly because the *No* coordination strategy's performance deteriorates with more decommitments (this is especially clear in case with no adverse effects when the buyer is least careful), but also because the performance of *Risk-Averse* strategy slightly improves with more and more possibilities for finding a very good deal.[2] This is consistent with hypothesis 31 so we can accept it.

An interesting additional observation is that the performance under the *No* coordination strategy behaves oddly when the effect is 0.1. Unlike in all other cases, the performance drops sharply when the first few markets are added, but from the fifth market upwards the performance actually seems to improve, although it never achieves the level with only one market. This type of peculiar behaviour is to be expected when there is no coordination because the combined effect of several markets may sometimes be surprising. In this instance, the reason for this behaviour is that the number of decommitments increases sharply with two markets (from zero to 9.18) and from there it keeps increasing until with four markets it reaches a maximum of 15.42, after which it decreases. This explains the changes in the performance.

The reason for this behaviour, on the other hand, is relatively simple. Because of the future offers, the *Adaptive Counter* often first finds contracts with opponents using the *Exponential Time-Dependent* tactic and those negotiations end around the same time (close to the end of negotiations) and it seems that this occurs often after the same matching, so when the number of markets increases so do the simultaneous extra contracts. However, there is a small chance that a good enough contract could be found in the previous matching or that good enough opponents using the *Random* tactic are present, both of which decrease the number of extra contracts. In case of earlier contracts, because fewer contracts are made that early and, in case of the *Random* tactic because the probability of simultaneous contracts are smaller (and because in those cases they will get into a contract before the opponents using the *Exponential Time-Dependent* tactics

[2] Given that the maximum theoretical performance in the setting is 50 and the actual average performance with 10 markets with no adverse effects is 49.82.

do). So after a few markets, when the number of markets increases, the number of decommitments decreases and the performance improves. Usually, of course, the number of decommitments increases when the number of markets increases, but as this example shows, this is not always the case.

This type of unexpected behaviour can occur when no coordination is used because there may be patterns in the service markets which are not visible with one market, but become visible when more markets are used. That is one more reason to use effective coordination, which is less susceptible to such effects because it tries to keep them under control.

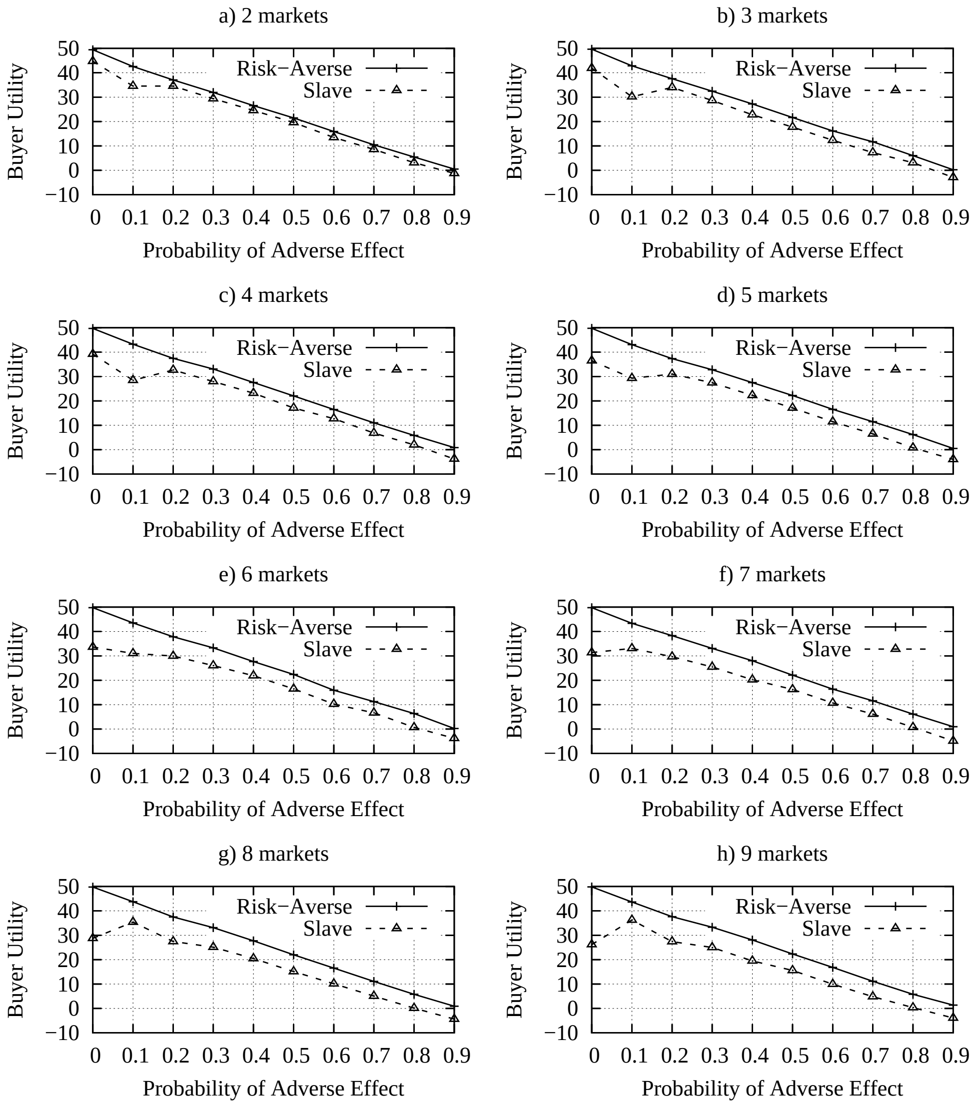

FIGURE 11.2: Risk-Averse and Slave coordination Strategies in the *Substitutes* Setting (Hypothesis 31).

We now move to consider the *Complements* setting. The relative performances of the *All or Nothing* and *No* coordination strategies are in figure 11.3 in the case where there are no market failures. The results have three different cases. First, with very low effect probabilities, coordination is counterproductive at least in the form of the *All or Nothing* strategy. The *No* coordination strategy beats the *All or Nothing* strategy in all cases when there are no adverse effects (at the $p < 0.0001$ level), in all cases except 2 markets cases when the effect probability is 0.1 (at the $p < 0.0001$ level except for 3 and 4 markets cases at the $p < 0.01$ level) and in the 7 markets case when the effect probability is 0.2. This is because the *Adaptive Counter* tactic is usually able to find a contract with a very good probability at some point, so it fails only rarely to get all the necessary contracts. On the other hand, the *All or Nothing* strategy is very conservative and only negotiates when it can be confident it can attain a contract in all markets at the same time. This means that with low effect probabilities the *All or Nothing* may wait and end up with inferior contracts because the best options may not be available at the same time.

The second part of the results is that with the moderate effect probabilities there is no statistically significant difference between the two. This means that the edge the coordination had with low effect probabilities is no longer there. This is because the higher probability of an adverse impact means that the best strategy in all markets is to wait. This means that all contracts are formed closer to the deadline, which means that the *No* coordination strategy has to take whatever is available then.

The third observation is that the *All or Nothing* strategy is able to beat the *No* coordination cases with very high effect probabilities with 3 or more markets (probability 0.9, at the $p < 0.0001$ level except for 3 markets case at the $p < 0.001$ level). This is because the very high effect probability means that all negotiators wait until the last possible minute and this may mean sometimes they are unsuccessful. And when they are unsuccessful of course, the *All or Nothing* strategy may have many contracts to decommit from, whereas the *All or Nothing strategy* does not even negotiate in cases where the risk of that happening is too high. This was what we said might happen in hypothesis 32, so we can accept it.

However, when we introduce (more) market failures (situations where getting a contract is not possible for any reason), the *All or Nothing* strategy's prudence is vindicated. In figure 11.2.3, we see what happens when 10% of the time one of the markets fails to produce any result. This of course is devastating for the

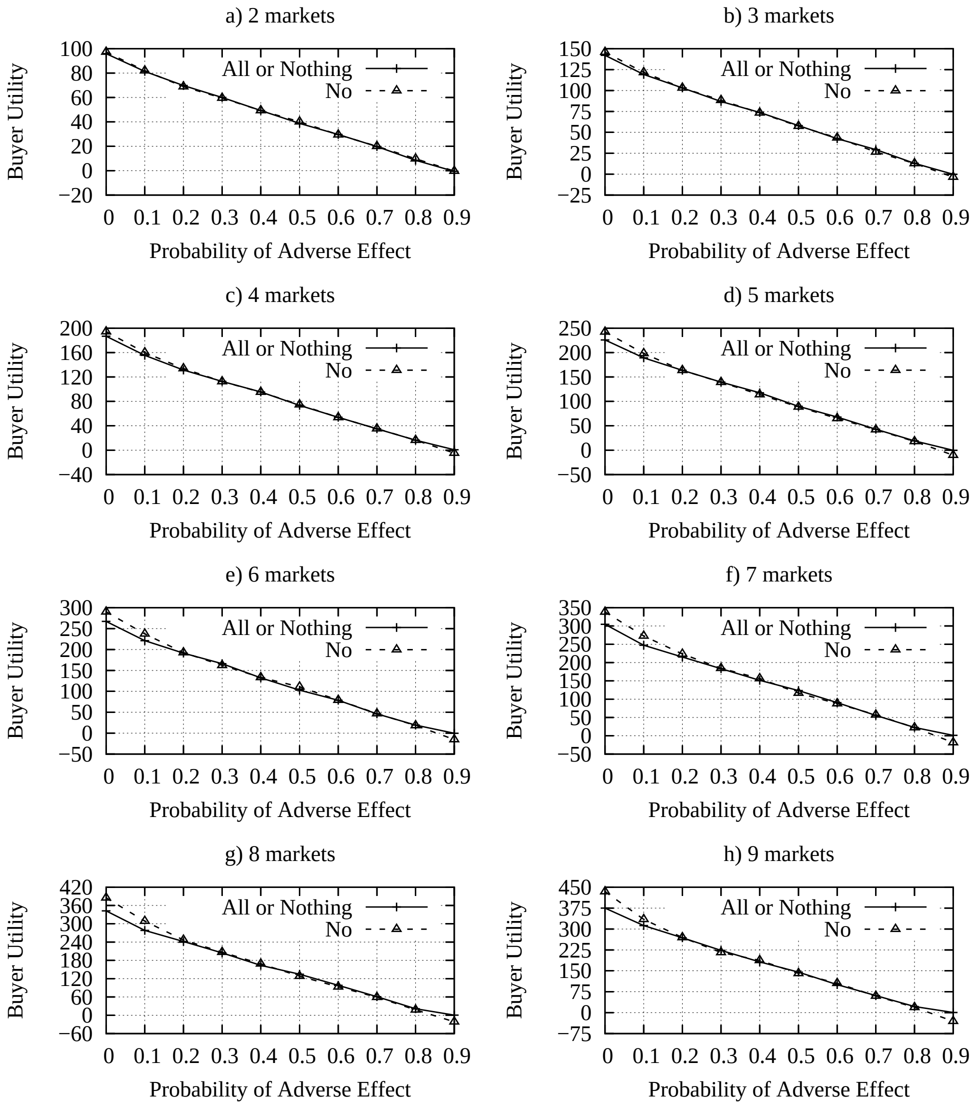

FIGURE 11.3: The performance of different coordination strategies in the *Complements* setting with no market failures (Hypothesis 32).

No coordination strategy case because it will often be able to get a contract in most unaffected markets and therefore in case of market failure, it will have to decommit from all the other contracts, which is very expensive. In contrast, the *All or Nothing* strategy does not enter into any contracts in case of market failures, but is able to get good results when there are no such events. It will therefore beat the *No* coordination strategy cases clearly in all cases (at the $p < 0.0001$ level). This is consistent with hypothesis 33 so we accept it.

In our setting, a 10% failure rate was enough, but in other settings more failures might be needed (at least in some cases). This is because here the buyer succeeded in each market with a very high probability if there was no market failure, which

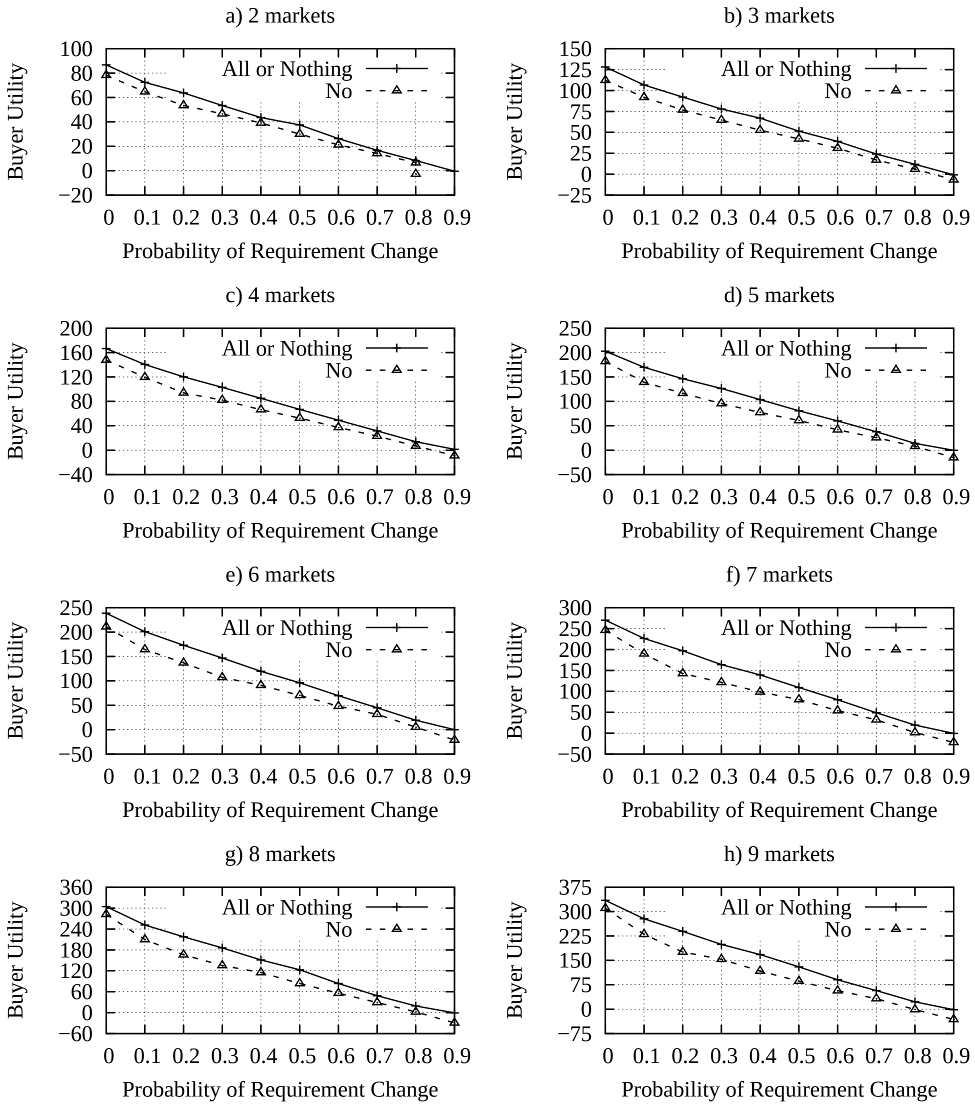

FIGURE 11.4: The performance of different coordination strategies in the *Complements* setting with 10% market failures (Hypothesis 33).

means that in case of market failure, the *No* coordination strategy had to pay very high decommitment fees. Now, if succeeding was less certain, the *No* coordination strategy would have to pay less fees in case of market failure. Moreover, if the window of opportunity in different markets would be at different times, the *No* coordination strategy would be successful more often when it is possible, whereas the *All or Nothing* strategy would rarely negotiate at all. In such environments, the *All or Nothing* strategy might need a market failure rate of more than 10% to be beneficial, especially with a high number of markets.

11.3 Summary

In this chapter, we have introduced the third and final level of our adaptive concurrent bilateral negotiation model. We showed how it can improve the performance of the buyer agent when there are interrelated services (substitutes or complements) it negotiates concurrently on (contribution **C8**). The coordination strategies we used were reasonably straight-forward and conservative and more sophisticated versions would undoubtedly work even better.

We also identified some environmental characteristics that make coordination especially useful. With the substitutes, the number of possible substitutes (market size) is essential. The performance of the *Slave* strategy deteriorates significantly and the performance of the *Risk-Averse* strategy improves slightly when the number of market increases. However, both, but especially the *Slave* strategy, would probably benefit from using future offers that take into account the number of markets instead of just one market as we had in the experiments. Some performance difference would be likely to remain, however, since the future offers are not an exact tool and it would not remove the need to decommit entirely when no coordination is used. In the complements case, we discovered that if a solution can be achieved in all markets with very high probability, the coordination strategy is probably not going to be that helpful. However, if there is a chance of failure in some markets, the coordination quickly becomes essential.

Chapter 12

Conclusions and Future Work

We have now detailed all the work in this thesis and now it is time to draw conclusions on our models and results. We do that in section 12.1. After that we will discuss some possible future directions for this work in section 12.2.

12.1 Conclusions

Generally, we can say that our work has shown that decommitment policies play a significant role in the welfare of all parties in the marketplace. During any opponent selection, negotiation and after the contract has been entered into, a rational party makes several decisions where the decommitment policy plays a significant role and which can have a significant impact on both his and his opponents' utilities. The effect of decommitment policies is not straight-forward and not always clear. This is because they can have very different effects on different decisions made by different parties at different times during the process. A policy that is optimal in terms of one decision, may be catastrophic in terms of another. On the other hand, when both parties take these decisions into account, they may find mutually acceptable solutions under most decommitment policies, so the role of decommitment policies should not be overstated. The decommitment policies do play a role, but it is the decisions made by the parties that are important.

Our work was divided into two parts, so we follow the same division in the conclusions as well. We first discuss our results in commitment models (section 12.1.1) and then our model and results in concurrent bilateral negotiation strategies (section 12.1.2).

12.1.1 Commitment Models

In the market setting, we investigated the common good and the effect the four basic decisions — performance, reliance, contract and selection — the parties make will have on it. Our aim was to investigate these decisions in a dynamic service market and see how taking these decisions into account changes the behaviour of the parties and how that changed behaviour changes the common good (see section 1.3 for what we set out to do). Much of the work was based on simple models and principles from law and law and economics and our contribution was to apply them to a dynamic service market setting and draw from them to make new decommitment policies and see how they fare against some known policies. We were also interested in settings with incomplete information and how well the sub-optimal policies worked.

In more detail, with the performance decision (chapter 4, contribution **C1**), we considered the effect of one or two-sided potential decommitments, re-entries and incomplete information. Now, for most of these settings the performance decision has not been previously discussed in such detail in the literature. We discovered that the best approach is to compensate for the victim's actual loss or under incomplete information, the best approximation for that loss. This proved very effective in most settings.

With the reliance decision (chapter 5, contribution **C2**), we considered a case where the buyer is able to enhance its utility by making investments in anticipation of the seller's performance and the reliance (performance probability) of the sellers vary. When the seller performs, the buyer's reliance on the performance has been beneficial, but if the seller does not perform, the buyer's reliance has been detrimental to himself and the common good because the cost cannot be retrieved. However, we discovered (as the models in the literature suggested) that taking the reliance decision into account improves the buyer's utility over some simple reliance strategies and any compensation for this reliance leads to over-reliance which can be detrimental to the common good.

In relation to the contract decision (chapter 6, contribution **C3**), we investigated the possibility that the parties use the contract price as a risk-allocation tool that allows them to find mutually acceptable (and beneficial to the society) contract prices in different situations and under different decommitment policies. We investigated the effect of none, one or both parties taking possible adverse effects into account and found out that for the common good, the best approach was that both parties took the contract decision into account, although in some situations it was enough that one of the parties did it.[1] Also parties were able to get similar utilities under most decommitment policies. There were some exceptions with the *Constant 1.0* policy, which occationally made it impossible for the parties to find a good deal.

With the selection decision (chapter 7, contribution **C4**), we investigated a setting where the buyers got to choose which sellers they wanted to negotiate with. Here, we discovered that compensatory policies made the buyers indifferent between performances and non-performances, although of course only the performance could produce utility for the society. Moreover, over-compensatory policies were even worse, because they made the buyers choose the least reliable providers. The best approach here was to have no decommitment fee at all or even have a decommitment bonus (the buyer would have to pay the seller in case the seller decommits) instead.

Although we discussed only a limited number of settings, we believe that the results mentioned above are generalizable to many other settings including those with different number of agents. Obviously, the number of agents does have a direct impact on the number of contracts created and, therefore, on the total utility gained, but the differences are likely to remain. However, because of the high variance of the results in a single contract, a reasonable number of contracts (agents) is often needed to see differences that are statistically significant. This means that the differences between different policies might, in some cases, be drowned in the noise, if the number of contracts (agents) is much lower than it was in our experiments.[2]

As already discussed in the literature review (section 2.2.3), the problem with the different decisions is that the optimal policy they prescribe varies and the optimal policy may be difficult to find in a setting where the different decisions are made in relation to the same negotiation or contract. For example, the selection decision suggests setting the seller's decommitment fee to zero or even negative, whereas such policy does not really work well with the performance decision. This is because the basic approach — which is to internalise the profits and costs of the opponent in the decision-maker's utility function — works very well only if there

[1]On the others, one party taking the decision into account could sometimes lead to worse outcomes than if neither had considered the contract decision.

[2]Moreover, we of course varied the number of sellers in the selection decision and showed that there the opponent selection becomes increasingly important when the number of sellers increases. Also the numbers of highly reliable sellers had a clear impact on the results with the reliance decision.

is one decision-maker and preferrably one decision to make. The approach is more difficult to use, when there are many decisions and many decision-makers and these decisions are heavily interconnected. In a realistic electronic market setting, however, such complexities are more than likely.

However, our results in relation to the contract decision, suggest that the decommitment policy may not always be essential to the common good, but if all parties took all these decisions into account simultaneously, they might be able to find a contract price that would distribute the risks in a way that is acceptable to everybody. This would safeguard at least positive expected utilities for everybody (and the society) and with everybody guarding their own interest the outcomes would also be good for the society. This is because if there was ever a better deal than the one the parties have found, one party would always be willing to move to that better deal and would be able to give the other party significant incentives to want the move too (by, for example, lowering the price). The role of decommitment policies would be limited to facilitating this price-setting mechanism or, more accurately, avoiding getting in its way (as happened with the *Constant 1.0* policy in some cases in chapter 7).

However, taking all the relevant decisions into account at the same time is far from simple in a non-trivial setting. The parties have to consider the performance and the reliance decisions when they are considering the contract decision and all these when they are choosing the opponents to negotiate with (the selection decision). On the other hand, most relevant parameters are not fixed, but they can also be changed. For example, in our settings, reliability of the opponent was known and fixed. However, a party can influence its reliability by investing in good technology and redundancies, for example, by having backup hardware in case one regularly used fails. A high decommitment fee will make a party invest in his reliability, but a low fee may not. When you add incomplete information, reputation and other relevant factors into this, you have a very difficult problem that may not have an easy solution.

Against this background, our goal in this work has been to introduce these decisions to people working on agents and dynamic service markets and show that the decisions are important in any environment where negotiation, contracts and decommitments are present and that they should be considered by both agent and market designers. We have discussed these decisions and some factors involved in great detail in order to increase understanding on decommitments, decommitment policies and their role in dynamic service markets.

12.1.2 Concurrent Bilateral Negotiation Strategies

In the concurrent bilateral negotiation setting, we considered a case with one buyer agent and several seller agents and with the buyer agent possibly negotiating with many sellers at the same time. Here, we considered two of the four decisions in the first part (contract and selection decision) in a complicated setting with sellers that do not consider any decisions. We improved and extended the state of the art in several ways (see section 1.3 for what we set out to do).

First, our adaptive concurrent bilateral negotiation model (chapter 8, contribution **C5**) is the most advanced model of its kind. It isolates the different relevant questions to specialised modules that can be freely changed to offer different kinds of functionality. The interaction between these modules is clearly specified and can facilitate very different kinds of behaviour and the different levels of the model offer guidance as to what sort of questions should be managed at each level. Some of the modules manage issues that are not discussed in the earlier literature, such as opponent selection, concurrency control and managing interrelated negotiations. This new functionality is offered in a very clear and easy-to-follow structure. The model is also able to accommodate taking the future offers (offers in later negotiations) into account and possibilities of adverse changes on either the buyer's or the seller's side.

In the lowest, `Negotiator`, level (chapter 9, contributions **C6** and **C9** (in part)) we allowed the sellers to use negotiation tactics that are based on either randomness or behavioural input from the buyer's part, which made it more difficult for the buyer to exploit his opponents even with full information. We introduced an optimal counter-tactic to each of the four seller tactics we allowed and a combination tactic, *Adaptive Counter*, that can adapt to any of the four tactics we allowed the sellers to have. Our counter tactics were designed to be used against a large group of opponent tactics and instead of trying to figure out what tactic the opponent is using, they use the seller's fundamental properties (such as reservation price) and an estimate on the distribution of the lowest offer to determine a single offer that will be made repeatedly in a negotiation. By so doing they make the outcome of a successful negotiation easier to estimate.

We found out that while the *Adaptive Counter* tactic was able to outperform the simpler countertactics with future offers and full information, it was much less effective with less information. In particular, it turned out that countertactics that simply assumed the opponent was using a certain tactic and made an optimal offer using that assumption did rather well (also with less information), when they had a reasonable chance of finding somebody who actually did use that tactic (or at least was willing to accept the offer made). This type of tactic might be able to obtain reasonable results in any market, no matter what tactics the opponents used. It only needs an estimate of how likely it is that certain kind of offers will be accepted in the marketplace given the number of opponents it will be able to encounter.

We made our biggest contributions in the middle level, where the `Controller`s operate (chapter 10, contributions **C7** and **C9** (in part)). Here, we discussed and empirically validated our ideas on opponent selection and concurrency control (choosing the number of opponents to have). In many settings, using more advanced opponent selection methods and concurrency strategies improved the outcome, although often the move from the *Quality* to the *Expected Utility* opponent selection proved to be less significant than the move from the *Random* to the *Quality* opponent selection. The *Analytic* and especially the *Adaptive* concurrency strategies were able to beat all the *Simple* strategies in most settings and with future offers, the *Simple* strategies were only better in a few cases where negotiating with as many opponents as possible was useful (the *Random* opponent selection with a very low probability of adverse effect) and the risks of extra contracts were neglible (the *Constant 0.0* decommitment policy). Also considering offers in later negotiations (future offers) turned out to be a very beneficial approach in almost all of our settings. Also here, the *Adaptive Counter* tactic often produced slightly better results than the *Random Counter* tactic under complete tactic information, but when the tactic information was completely unavailable, the performance of *Adaptive Counter* took a significant hit, whereas with the *Random Counter* the effect was much smaller (and only existing at all with the *Analytic* concurrency strategy). In other words, we showed that the concurrency strategies, opponent selection and considering future offers were all essential to the concurrent bilateral negotiation.

We also discussed managing interconnected negotiations at the `Coordinator` level (chapter 11, contribution **C8**). Here, we showed that even very basic coordination strategies can improve the performance of the buyer agent with both substitutes and complements.

In sum, we believe our adaptive concurrent bilateral negotiation model is very effective and shows how different issues can be effectively separated from each other and how simple strategies can be used to improve the buyer's utility. We believe it can provide a good starting point for even more sophisticated concurrent negotiation agents since it is easily extensible.

As in Part I, here the results are generalizable. We only have one buyer and, therefore, at most one contract after each run and we had to compare groups of 100 runs to show differences. In a single run or in a group of a few runs, the differences between the strategies would often be drowned in the noise. We could vary the number of sellers, however, without much change in our results. The more advanced forms of opponent selection (quality, expected utility) do better when there are more seller agents in the market because that usually means that the best agents they can find are better than with fewer agents (the best agent out of 1000 agents is probably going to be better than the best agent out of 10 agents if the agents are drawn from the same distribution). But at least the advanced strategies would usually find a contract also with a very few sellers and the differences in performance to a setting with a large number of sellers would probably be relatively small.

12.2 Future Work

Finally, we will discuss possible directions in which our work could be taken. Again, we will discuss the two parts separately: commitment models in section 12.2.1 and concurrent bilateral negotiation strategies in section 12.2.2.

12.2.1 Commitment Models

In the context of Part I, the obvious directions of future work would be to:

- consider more than one decision simultaneously. The easiest combination would be to consider the performance and/or reliance decision with the contract decision. Considering these decisions already in the selection decision would be more complex, as well as adding other decisions such as preparation decision (where a party can influence its reliability by investing into it). The risk here is that any setting with more than one decision would be quite complex and possibly too specific to make any generalisations that would apply to other settings. It would be, nonetheless, interesting to see how the interplay between different decisions works and how the parties would be able to use the contract price as a risk-allocation tool in a more complex setting.

- allow both parties to consider the decision(s) at the same time. In all but the contract decision, we did not allow more than one of the parties to consider its position. For example, in the selection decision, a high decommitment fee would not only mean that the buyers would try to choose as unreliable sellers as possible (because they prefer non-performance), but it would also mean that the unreliable providers would ask for a very high contract price to justify the high risk of non-performance and high fee or, if that was not possible, they would not enter into market at all. In such a setting, a high fee might therefore mean self-selection of the sellers leaving only high reliability providers left. In situations where high reliability is difficult (expensive) to attain, a high decommitment fee could mean that these services are not provided at all.

Another possible approach would be to try to make the market setting somewhat more realistic. This could, for example, involve using future offers (from part II) so that parties would also consider future matchings when they are considering if they want to enter a contract or not. And, of course, in a realistic setting all parties would not be using the *Exponential Time-Dependent* tactic with β parameters selected at random. Here, we could also remove the limitation that our agents will not engage in strategic behaviour in their negotiation, meaning that they will accept any acceptable deal and do not try to outwait each other, for example. Although this was not critical to our work (because we aimed to investigate these decisions and such considerations would have distracted us from our task), we would have to consider such behaviour in a more realistic setting. On the other hand, we might end up with something like what we had in the concurrent negotiation part, because completely strategic behaviour would require full game theoretic approach and, as we discussed in section 2.2.1.1, this would be a very difficult road to take.

12.2.2 Concurrent Bilateral Negotiation Strategies

In the concurrent bilateral negotiation, the most interesting directions for future research include:

- extending the approach to multi-attribute negotiations. In section 2.3.1.4, we suspected that the concurrent bilateral negotiation could truly be in multi-attribute settings (especially as compared to auctions). We also said that we investigated concurrent bilateral negotiation in price negotiations only as a step towards the multi-attribute settings. In the future work, it would therefore be important to take the necessary steps to facilitate using concurrent bilateral negotiation in multi-attribute settings, for example, configuring the services to the buyer's needs.

- investigating settings with more opponent tactics. Our work here had only four negotiation tactics for the seller, but in any realistic free marketplace there could be any number of them. An interesting approach would be to investigate if we could use something similar to our countertactics even if we do not know anything about the tactics the opponents use, but can collect information about their offers freely. In theory, over time, this should allow us to estimate probabilities of success for a given offer even if we knew very little about the opponents and their negotiation tactics. This is especially true over large number of negotiations. This approach could give us reasonable results in any opponent pool and with very limited starting information. In particular, it would be interesting to investigate machine learning techniques for getting and updating these distributions on the go.

- investigating settings with less information. For example, the advanced concurrency strategies are based on a decision-theoretic approach which means that they need quite a lot of reasonably accurate information to be useful. Now, especially problematic is the fact the distribution of contract times, which may not be available if the opponent tactics are simply not known or even restricted in any way. It might therefore be useful to use empirical data or some heuristics instead. Also here it would be interesting to investigate machine-learning techniques and start with no information. And we could also have uncertainty about the reservation prices. Although we assumed that the parties know each others' reservation price exactly, it is only used as a way to estimate the probability distribution for their lowest offer. Therefore, incorporating uncertain reservation prices is quite straight-forward. On the other hand, if we can have that distribution of the lowest offer by other means, we do not need the reservation price at all.

- extending the coordinator level to manage more complex interrelations between contracts. In particular, another area where the concurrent negotiation could be useful (see our discussion in section 2.3.1.4) would be in managing interconnected negotiations. The first step might be introducing something similar to the *Analytic* and/or *Adaptive* concurrency strategies at the `Coordinator` level and also extending future offers to take into account the possibilities a greater selection would offer in the substitutes case. As we speculated in section 11.1, the next steps might have to do with more complex settings where both complements and substitutes would be combined in a single case.

Bibliography

Admati, A. R. and Perry, M. 1987. Strategic delay in bargaining. *Review of Economic Studies*, 54(3):345–364.

Ancarani, F. 2002. Pricing and the Internet: Frictionless commerce or pricer's paradise? *European Management Journal*, 20(6):680–687.

Andersson, M. R. and Sandholm, T. W. 2001. Leveled commitment contracts with myopic and strategic agents. *Journal of Economic Dynamics & Control*, 25:615–640.

Atiyah, P. S. 1995. *An Introduction to the Law of Contract*. Oxford University Press, 5th edition.

Atiyah, P. S., Adams, J. N., and MacQueen, H. 2005. *Sale of Goods*. Longman, 11th edition.

Aumann, R. J. 1997. Rationality and bounded rationality. *Games and Economic Behavior*, 21:2–14.

Ausubel, L. M. and Deneckere, R. J. 1992. Bargaining and the right to remain silent. *Econometrica*, 60(3):597–625.

Bajari, P., McMillan, R., and Tadelis, S. 2008. Auctions versus negotiations in procurement: An empirical analysis. To appear on *Journal of Law, Economics and Organization*, advance access available at http://jleo.oxfordjournals.org/cgi/content/abstract/ewn002.

Barton, J. H. 1972. The economic basis of damages for breach of contract. *Journal of Legal Studies*, 1(2):277–304.

Berners-Lee, T., Hendler, J., and Lassila, O. 2001. The semantic web. *Scientific American*, 284(5):34–43.

Binmore, K. G. and Herrero, M. 1988. Matching and bargaining in dynamic models. *Review of Economic Studies*, 55:17–31.

Brans, J. P. and Vincke, P. 1985. A preference ranking organisation method (the promethee method for multiple criteria decision-making). *Management Science*, 31(6):647–656.

Bui, H. H., Venkatesh, S., and Kieronska, D. 1999. Learning other agents' preferences in multi-agent negotiation using the bayesian classifier. *International Journal of Cooperative Information Systems*, 8(4):275–294.

Chajewska, U., Koller, D., and Ormoneit, D. 2001. Learning and agent's utility function by observing behavior. In *Proceedings of the Eighteenth International Conference on Machine Learning (ICML'01)*, pages 35–42, Williamstown, MA, USA.

Chatterjee, K. and Dutta, B. 1998. Rubinstein auctions on competition for bargaining partners. *Games and Economic Behavior*, 23:119–145.

Chatterjee, K. and Samuelson, L. 1987. Bargaining with two-sided incomplete information: An infinite horizon model with alternating offers. *The Review of Economic Studies*, 54(2):175–192.

Cho, I.-K. 1990. Uncertainty and delay in bargaining. *The Review of Economic Studies*, 57(4):575–595.

Coehoorn, R. M. and Jennings, N. R. 2004. Learning an opponent's preferences to make effective multi-issue negotiation tradeoffs. In *Proceedings of the 6th International Conference on E-Commerce (ICEC'04)*, pages 59–68, Delft, The Netherlands.

Collins, J., Bilot, C., Gini, M., and Mobasher, B. 2001. Decision processes in agent-based automated contracting. *IEEE Internet Computing*, pages 61–72.

Cramton, P. C. 1984. Bargaining with incomplete information: An infinite-horizon model with two-sided uncertainty. *Review of Economic Studies*, 51:579–593.

Cramton, P. C. 1992. Strategic delay in bargaining with two-sided uncertainty. *Review of Economic Studies*, 59:205–225.

Craswell, R. 1988. Contract remedies, renegotiation, and the theory of efficient breach. *South California Law Review*, 61(3):629–670.

Craswell, R. 2001. The two economic theories of enforcing promises. In Benson, P., editor, *The Theory of Contract Law*, pages 19–44. Cambridge University Press.

Curbera, F., Duftler, M., Khalaf, R., Nagy, W., Mukhi, N., and Weerawarana, S. 2002. Unraveling the Web services web: An introduction to SOAP, WSDL, and UDDI. *IEEE Internet Computing*, 6(2):86–93.

Das, R., Hanson, J. E., Kephart, J. O., and Tesauro, G. 2001. Agent-human interactions in the continuous double auction. In *Proceedings of the International Joint Conferences on Artificial Intelligence (IJCAI) 2001*, pages 1169–1176.

David, E., Rogers, A., Schiff, J., Kraus, S., and Jennings, N. R. 2005. Optimal design of english auctions with discrete bid levels. In *Proceedings of 6th ACM Conference on Electronic Commerce (EC'05)*, pages 98–107.

Excelente-Toledo, C. B., Bourne, R. A., and Jennings, N. R. 2001. Reasoning about commitments and penalties for coordination between autonomous agents. In *Proceedings of 5th International Conference on Autonomous Agents (AGENTS'01)*, pages 131–138, Montreal, Canada.

Faratin, P., Sierra, C., and Jennings, N. R. 1998. Negotiation decision functions for autonomous agents. *International Journal of Robotics and Autonomous Systems*, 24(3-4):159–182.

Fatima, S. F., Woodridge, M., and Jennings, N. R. 2004. An agenda-based framework for multi-issue negotiation. *Artificial Intelligence*, 152(1):1–45.

Ferguson, T. S. 2008. Optimal stopping and applications. Available at http://www.math.ucla.edu/∼tom/Stopping/Contents.html.

Foster, I. and Kesselman, C., editors 2004. *The Grid: Blueprint for a New Computing Infrastructure*. Morgan Kaufmann, 2nd edition.

Fudenberg, D. and Tirole, J. 1983. Sequential bargaining with incomplete information. *The Review of Economic Studies*, 50(2):221–247.

Fudenberg, D. and Tirole, J. 1991. *Game Theory*. MIT Press.

Fuller, L. and Purdue, W. 1936. The reliance interest in contract damages. *Yale Law Journal*, 46:52–96.

Goetz, C. J. and Scott, R. E. 1980. Enforcing promises: An examination of the basis of contract. *Yale Law Journal*, 89(7):1261–1288 and 1321–1322.

Grover, V. and Ramanlal, P. 1999. Six myths of information and markets: Information technology networks, electronic commerce, and the battle for consumer surplus. *MIS Quarterly*, 23(4):455–485.

He, M., Jennings, N. R., and Leung, H.-F. 2003. On agent-mediated electronic commerce. *IEEE Transactions on Knowledge and Data Engineering*, 15(4):985–1003.

Hou, C. 2004. Predicting agent tactics in automated negotiation. In *Proceedings of IEEE/WIC/ACM International Conference on Intelligent Agent Technology (IAT'04)*, pages 127–133, Beijing, China.

Huhns, M. N. 2002. Agents as web services. *IEEE Internet Computing*, 6(4):93–95.

Huhns, M. N. and Singh, M. P. 2005. Service-oriented computing: Key concepts and principles. *IEEE Internet Computing*, 9(1):75–81.

Janssen, M. and Rasmusen, E. 2002. Bertrand competition under uncertainty. *The Journal of Industrial Economics*, 50(1):11–21.

Jennings, N. R. 1993. Commitments and conventions: The foundation of coordination in multi-agent systems. *The Knowledge Engineering Review*, 8(3):223–250.

Jennings, N. R. 2001. An agent-based approach for building complex software systems. *Communications of the ACM*, 44(4):35–41.

Jennings, N. R., Faratin, P., Lomuscio, A. R., Parsons, S., Sierra, C., and Wooldridge, M. 2001. Automated negotiation: Prospects, methods and challenges. *International Journal of Group Decision and Negotiation*, 10(2):199–215.

Kambil, A. and van Heck, E. 2002. *Making Markets*. Harvard Business School Press.

Keeney, R. L. and Raiffa, H. 1976. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. John Wiley & Sons.

Kephart, J. O., Hanson, J. E., and Greenwald, A. R. 2000. Dynamic pricing by software agents. *Computer Networks*, 32:731–752.

Kornhauser, L. A. 1986. An introduction to the economic analysis of contract remedies. *University of Colorado Law Journal*, 57:683–725.

Krishna, V. 2002. *Auction Theory*. Academic Press, Elsevier Science.

Lehmann, D., Müller, R., and Sandholm, T. 2006. The winner determination problem. In Smith, V. L., Cramton, P., Shoham, Y., and Steinberg, R., editors, *Combinatorial Auctions*, chapter 12, pages 297–318. MIT Press.

Li, C., Giampapa, J., and Sycara, K. 2004. Bilateral negotiation decisions with uncertain dynamic outside options. In *Proceedings of the First IEEE International Workshop on Electronic Contracting (WEC'04)*, pages 54–61, San Diego, CA, USA.

Li, C., Giampapa, J., and Sycara, K. 2006. Bilateral negotiation decisions with uncertain dynamic outside options. *IEEE Transactions on Systems, Man, and Cybernetics, Part C*, 36(1):3–10.

Li, C., Sycara, K., and Giampapa, J. 2005. Dynamic outside options in alternating-offers negotiations. In *Proceedings of 38th Annual Hawaii International Conference on System Sciences (HICSS'05)*, Waikoloa, HI, USA.

Lin, M., Xie, J., Guo, H., and Wang, H. 2005. Solving qos-driven web service dynamic composition as fuzzy constraint satisfaction. In *Proceedings of the IEEE International Conference on e-Technology, e-Commerce and e-Service (EEE 2005), Hong Kong, China*, pages 9–14.

Lomuscio, A. R., Wooldridge, M., and Jennings, N. R. 2003. A classficiation scheme for negotiation in electronic commerce. *International Journal of Group Decision and Negotiation*, 12(1):31–56.

Mas-Colell, A., Whinston, M. D., and Green, J. R. 1995. *Microeconomics Theory.* Oxford University Press.

Mok, W. W. H. and Sundarraj, R. P. 2005. Learning algorithms for single-instance electronic negotiations using the time-dependent behavioral tactic. *ACM Transactions on Internet Technology*, 5(1):195–230.

Muthoo, A. 2002. *Bargaining Theory and Applications.* Cambridge University Press.

Nguyen, T. D. and Jennings, N. R. 2003. A heuristic model for concurrent bi-lateral negotiations in incomplete information settings. In *Proceedings of 18th International Joint Conference on Artificial Intelligence (IJCAI'03)*, pages 1467–1469, Acapulco, Mexico.

Nguyen, T. D. and Jennings, N. R. 2004. Coordinating multiple concurrent negotiations. In *Proceedings of 3rd International Conference on Autonomous Agents* *and Multi-Agent Systems (AAMAS'04)*, pages 1064–1071, New York City, NY, USA.

Nguyen, T. D. and Jennings, N. R. 2005. Managing commitments in multiple concurrent negotiations. *International Journal of Electronic Commerce Research and Applications*, 4(4):362–376.

Osborne, M. and Rubinstein, A. 1990. *Bargaining and Markets*. Academic Press.

Perry, M. 1986. An example of price formation in bilateral situations: A bargaining model with incomplete information. *Econometrica*, 54(2):313–322.

Pine, B. J. 1993. *Mass Customization: The New Frontier in Business Competition*. Harvard Business School Press.

Ponka, I. and Jennings, N. R. 2007. Changing circumstances and leveled commitment: A compensatory approach to contracting. In *Proceedings of the IEEE Joint Conference on E-Commerce Technology (CEC) and Enterprise Computing, E-Commerce and E-Services (EEE) (CEC-EEE 2007)*, pages 141–148.

Rahwan, I., Kowalczyk, R., and Pham, H. H. 2002. Intelligent agents for automated one-to-many e-commerce negotiation. In Ousdhoorn, M., editor, *Proceedings of 25th Australian Computer Science Conference (ACSC2002)*, pages 197–204, Melbourne, Australia.

Rahwan, I., Ramchum, S. D., Jennings, N. R., McBurney, P., Parsons, S., and Sonenberg, L. 2003. Argumentation-based negotiation. *The Knowledge Engineering Review*, 18(4):343–375.

Raiffa, H. 1987. *The Art and Science of Negotiation*. Belknap Press, reprint edition.

Rosenschein, J. S. and Zlotkin, G. 1994. *Rules of Encounter*. MIT Press.

Rothkopf, M. H. and Harstad, R. M. 1994. Modeling competitive bidding: A critical essay. *Management Science*, 40(3):364–384.

Rubinstein, A. 1982. Perfect equilibrium in a bargaining model. *Econometrica*, 50(1):97–109.

Rubinstein, A. and Wolinsky, A. 1985. Equilibrium in a market with sequential bargaining. *Econometrica*, 53(5):1133–1150.

Russell, S. and Norvig, P. 2003. *Artificial Intelligence: A Modern Approach.* Prentice Hall, 2nd edition.

Saaty, T. L. 1980. *The Analytical Hierarchy Process.* McGraw-Hill.

Salanié, B. 2000. *The Economics of Contracts: A Primer.* MIT Press.

Sandholm, T. and Lesser, V. 1996. Advantages of a leveled-commitment contracting protocol. In *Proceedings of the 13th National Conference on Artificial Intelligence (AAAI-96)*, pages 127–133, Portland, OR, USA.

Sandholm, T., Sikka, S., and Norden, S. 1999. Algorithms for optimizing leveled commitment contracts. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI'99)*, pages 535–540, Stockholm, Sweden.

Sandhom, T. and Lesser, V. 2001. Leveled commitment contracts and strategic breach. *Games and Economic Behavior*, 35:212–270.

Sandhom, T. and Lesser, V. 2002. Leveled commitment contracting: A backtracking instrument for multiagent systems. *AI Magazine*, 23(3):89–100.

Sen, S. and Weiss, G. 1999. Learning in multiagent systems. In Weiss, G., editor, *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*, pages 259–298. MIT Press.

Seo, Y.-J., Jeong, H.-Y., and Song, Y.-J. 2004. A study on web services selection method based on the negotiation through quality broker: A maut-based approach. In *Proceedings of International Conference on Embedded Software and Systems (ICESS 2004), Hangzhou, China*, pages 65–73.

Seo, Y.-J., Jeong, H.-Y., and Song, Y.-J. 2005. Best web service selection based on the decision making between qos criteria of service. In *Proceedings of the Second International Conference on Embedded Software and Systems (ICESS 2005), Xi'an, China*, pages 408–419.

Sim, K. M. 2002. A market-driven model for designing negotiation agents. *Computational Intelligence*, 18(4):618–637.

Sim, K. M. 2004. Negotiation agents that make prudent compromises and are slightly flexible in reaching consensus. *Computational Intelligence*, 20(4):643–662.

Sim, K. M. and Wang, S. Y. 2004. Flexible negotiation agent with relaxed decision rules. *IEEE Transactions on Systems, Man, and Cybernetics – Part B: Cybernetics*, 34(3):1602–1608.

Simon, H. 1955. A behavioral model of rational choice. *The Quarterly Journal of Economics*, 64:99–118.

Smith, S. A. 2004. *Contract Theory*. Oxford University Press.

Spulber, D. F. 1995. Bertrand competition when rivals' costs are unknown. *The Journal of Industrial Economics*, 43(1):1–11.

Stirling, W. C. 2003. *Satisficing Games and Decision Making: With Applications to Engineering and Computer Science*. Cambridge University Press.

Treitel, G. H. 2003. *The Law of Contract*. Sweet & Maxell, 11th edition.

Varian, H. R. 2003. *Intermediate Microeconomics: A Modern Approach*. W.W.Norton & Company, 6th edition.

Vulkan, N. 2003. *The Economics of E-Commerce*. Princeton University Press.

Weiss, R. M. and Mehrotra, A. K. 2001. Online dynamic pricing: Efficiency, equity and the future of e-commerce. *Virginia Journal of Law and Technology*, 6(2), available at: http://www.vjolt.net/vol6/issue2/v6i2-a11-Weiss.html.

Wooldridge, M. 2002. *An Introduction to Multiagent Systems*. John Wiley & Sons, Ltd.

Zeng, D. and Sycara, K. 1998. Bayesian learning in negotiation. *International Journal of Human-Computer Studies*, 48:125–141.

Zeng, L., Benatallah, B., Dumas, M., Kalagnanam, J., and Sheng, Q. Z. 2003. Quality driven web services composition. In *Proceedings of the 12th International World Wide Web Conference (WWW2003)*, pages 411–421.

Zhang, X., Lesser, V., and Abdallah, S. 2005. Efficient management of multi-linked negotiation based on a formalized model. *Autonomous Agents and Multi-Agent Systems*, 10:165–205.