

Modeling Preferences, Strategic Reasoning
and Collaboration in
Agent-Mediated Electronic Markets

Graduation committee:

Prof. dr. A.G.L. Romme (chair)

Prof. dr. Han La Poutré (promotor)

Core reading committee:

Prof. dr. Nicholas R. Jennings (Univ. of Southampton, UK)

Prof. dr. David C. Parkes (Harvard University, US)

Prof. dr. Ton G. de Kok (Technische Universiteit Eindhoven)

Prof. dr. Emile H. L. Aarts (TU Eindhoven & Philips Research)

Additional members:

Prof. dr. Catholijn Jonker (Technical University of Delft, NL)

Prof. dr. Cees Witteveen (Technical University of Delft, NL)



Centrum Wiskunde & Informatica



SIKS Dissertation Series No. 2009-19

The research reported in this thesis has been carried out at Centrum voor Wiskunde en Informatica (CWI), Amsterdam under the auspices of SIKS, the Dutch Research School for Information and Knowledge Systems.

ISBN: 978-90-386-1816-6

A catalogue record is available from the Eindhoven University of Technology Library

Published by Technische Universiteit Eindhoven, 2009

Printed on demand through Lulu.com

Copyright © 2009 by Valentin Robu. All rights reserved.

Cover illustrations:

Front: “A fractal model of the brain” - courtesy of Sven Geier.

Back: The Mandelbrot set fractal.

Modeling Preferences, Strategic Reasoning and Collaboration in Agent-Mediated Electronic Markets

PROEFSCHRIFT

ter verkrijging van de graad van doctor aan de
Technische Universiteit Eindhoven, op gezag van de
rector magnificus, prof.dr.ir. C.J. van Duijn, voor
een commissie aangewezen door het College voor
Promoties in het openbaar te verdedigen
op donderdag 2 juli 2009 om 16.00 uur

door

Valentin Robu

geboren te Cluj-Napoca, Roemenië

Dit proefschrift is goedgekeurd door de promotor:

prof.dr.ir. J.A. La Poutré

Acknowledgements

First and foremost, I'm deeply grateful to my advisor, Han La Poutré for having me as his PhD student and for his patient guidance all these years. Our discussions provided me with lots of insights and motivation (at times when it was lacking), and had a defining contribution in shaping my research interests and career. I had a great deal to learn from Han, not only about what it means to get solid research results, but also about academic life in general.

Also, a special “thank you” to Catholijn Jonker, who, as my Master thesis supervisor at the Free University Amsterdam introduced me to the world of multi-agent systems, and to automated negotiation in particular. Her enthusiasm and dedicated guidance were highly motivating in determining me to pursue a PhD after graduation.

I am also grateful to Nick Jennings and David Parkes, my external thesis committee members for the careful reading of my dissertation and their suggestions, which made the presentation of this thesis much stronger. Furthermore, I would like to thank professors Ton de Kok, Emile Aarts and Cees Witteveen for reading my thesis and their insightful comments. I also acknowledge professor Ioan Alfred Letia (TU Cluj-Napoca, Romania), who first recommended me to pursue a master degree in the Netherlands.

This thesis would not be what it is without the excellent collaborators and colleagues I had over the years. I will begin by my co-authors (and, at times, co-advisers) at CWI: Koye Somefun, Pieter Jan 't Hoen, Sander Bohte, Han Noot, Peter Bosman and Enrico Gerding (whom I had the pleasure of meeting again in Southampton). The breadth of their research interests and the ideas we exchanged (some of which materialized in publications) were an asset to my research career.

Also, many thanks to prof. Cees Witteveen and the group from Delft: Mathijs de Weerd, Tomas Klos (who was also my room-mate at CWI for a year), Renze Steenhuizen, Yingqian Zhang, Sicco Verwer and Tamas Mahr. The invitation to lecture there (part-time) for a year was a great learning experience for me, since I got to see first hand the challenges involved in teaching side of academic life. Also Koen Hindriks and Dmytro Tykhonov from Delft provided many interesting discussions.

Another defining moment in my PhD research was attendance of the Santa Fe Complex Systems Summer School. I am in debt to all teachers and fellow participants, but in particular to Harry Halpin (Edinburgh and MIT) and Hana Shepherd (Princeton). The excellent collaboration forged with them led to a very interesting and succesfull line of research on collaborative tagging, which resulted in a chapter of this thesis. Furthermore, I had the pleasure of meeting Harry several times afterwards (which included sharing a room in Dagstuhl),

and his enthusiasm and insights motivated me to learn a lot about web science.

I would also like to thank Takayuki Ito (Nagoya University and MIT), Minjie Zhang (Wollongong University, Australia), Shaheen Fatima (Liverpool), Tokuro Matsuo (Nagoya) and Hiromitsu Hattori (Kyoto Univ. - who also visited us at CWI). I had the pleasure of co-organising with them a succesful workshop on complex automated negotiations (this year, for the third year running), as well as coordinating on an AAMAS conference tutorial.

I am also grateful to the people from Almende BV, Rotterdam, who offered me a year-long internship during my master thesis days, of which I mention here: Peet van Tooren, Hans Abbink, Tamas Mahr, JP Larsen and Stefan Kroon. Besides improving considerably my programming skills, this internship gave me a great perspective on the challenges involved in deploying agent systems in practice, which I would never get from a purely academic career. The lively arguments with Peet also taught me a lot about the importance of keeping an open mind and defending the practical relevance of one's work. Also, I would like to thank Willem-Jan van Schijndel, from our project partner Vos Logistics, for his time in working with us to get a concrete application from our work on auction-based allocation.

During my time at CWI, I have been privileged to supervise two excellent master students: Sander van der Putten (Technology Management, Eindhoven) and Lonneke Mous (Econometrics, Univ. Rotterdam). Both had excellent results (including cum laude graduations and prizes for best thesis in the fields in the Netherlands). Sander helped us close the conceptual gap between the way AI researchers and logistic managers see the world. The work with Lonneke on using priced options to solve the exposure problem in sequential auctions definitely led to one of the most interesting directions of work to emerge from this thesis, one that I definitely intend to pursue further in my research work. Moreover, Lonneke is a great person to work with, and I remember fondly the travels through the narrow (and sunny) streets of Lisbon with Lonneke and Anke.

Which brings me to one of the most important groups of people I met during my PhD. I'm referring, of course, to the other members of the "gang of four" PhD students of Han: Anke Hutzenschreuter, Ivan Vermeulen and Mengxiao Wu (of which, in that order, the first two were kind enough to be my paranimfem, while the last two were my room-mates at CWI for the past 3 years). Our discussions covered everything from bargaining models and distributed patient scheduling in hospitals to food recipes and the virtues of wheat beer. Also, our culinary explorations of the delights of Amsterdam will be very fondly remembered.

Also, I'd like to thank all my friends in Amsterdam, starting with my Romanian friends: Tudor, Radu, Ghita, Vlad and Liviu, who have been a great help and support when moving to the Netherlands, and were always there to share a beer in Cafe Uilenstede (and move furniture, another constant in the life of a PhD student). Also, many thanks to my many other Dutch friends: Benno, Abel, Richard, Leopold, Jan, Onno, Ruggiero and Joost. You showed me that doing a PhD in Amsterdam can have its fair share of excitement, besides from the professional excitement of browsing the latest AAMAS conference proceedings.

Finally, my family who have helped me through this all along, especially my mother Ioana and my sister Flaminia. Their support is invaluable.

Contents

1	Introduction	1
1.1	Agents and electronic markets	1
1.2	Negotiation (bargaining) vs. auction protocols	2
1.3	Designing for strategic behaviour: market mechanisms vs. individual agent strategies	3
1.4	Modeling preferences and utilities in agent-mediated market settings	5
1.5	Emergence of collaboration and structure in multi-agent systems	6
1.6	Positioning of the contributions of this thesis	7
1.7	Modeling of combinatorial preferences (multi-issue or multi-item) in bilateral negotiations	8
1.7.1	Pareto-optimal outcomes in multi-issue negotiation	9
1.7.2	Modeling multi-attribute negotiation with incomplete preference information	11
1.7.3	Non-linear and combinatorial preferences in negotiation	11
1.7.4	Modeling multi-item negotiations over k-additive utility functions using utility graphs	13
1.7.5	Individual preferences and social influence	14
1.7.6	Learning the structure of utility graphs used in multi-item negotiation through collaborative filtering	15
1.8	Preferences under uncertainty and bidding in sequential auctions	16
1.8.1	Sequential auctions and the exposure problem	16
1.8.2	Designing sequential auction strategies for risk-averse bidders	17
1.8.3	Options mechanisms in sequential options	18
1.8.4	Using priced options to solve the exposure problem	19
1.9	Applications to transportation logistics	20
1.9.1	Auction-based allocation of transportation loads in multi-party transportation logistics	20
1.10	Preferences in social web communities and online markets	21
1.10.1	The complex dynamics of collaborative tagging systems	22
1.10.2	An empirical analysis of sponsored search markets	23
1.11	Structure of the thesis	23
1.12	Publications related to each chapter	26
1.12.1	Research work not included in the current thesis	28

I	Modeling combinatorial preferences in bilateral negotiation (bar-gaining)	31
2	An Agent Architecture for Cooperative Multi-Attribute Negotiation With In-complete Preference Information	33
2.1	Introduction	33
2.2	The multi-attribute negotiation model	34
2.2.1	Bid Utility Determination and Planning Component	35
2.2.2	The Attribute Planning Component	36
2.2.3	The Target Evaluation Planning Component	37
2.2.4	Estimation of Opponent's Parameters Component	39
2.2.5	Guess Coefficients Component	40
2.3	Implementation & Experimental validation	41
2.3.1	Experimental set-up	41
2.3.2	An example negotiation trace	42
2.3.3	Comparing traces from the same test set	45
2.3.4	Comparing results from all test sets	48
2.3.5	Human-computer experiments	49
2.4	Discussion	50
2.5	Conclusions	51
3	Modeling Complex Multi-Issue Negotiations Using Utility Graphs	53
3.1	Introduction	53
3.2	The negotiation setting	57
3.2.1	Net Utility functions of Buyer and Seller	57
3.2.2	Using gains from trade as efficiency criteria	58
3.2.3	Outline of the negotiation setting and protocol	58
3.2.4	Assumptions about buyer knowledge	60
3.2.5	Top-level negotiation algorithm used by the seller	60
3.3	Decomposable utility functions and their graphical representation	61
3.3.1	The k-additive utility form	62
3.3.2	Using graphs to model complex utility functions	63
3.4	Negotiation heuristics based on utility graphs	64
3.4.1	Selecting the best counter offer	65
3.4.2	Updating Sub-utility Functions	69
3.5	Constructing the structure of utility graphs using aggregate negotiation data	72
3.5.1	Collaborative filtering: brief introduction	72
3.5.2	Overview of our filtering and negotiation approach	73
3.5.3	Minimal super-graph for a class of buyers	74
3.5.4	Extracting information from concluded negotiation data	75
3.5.5	Computing the similarity matrices	76
3.5.6	Building the super-graph of buyer utilities	77
3.5.7	Minimization of expected loss in Gains from Trade as cut-off criteria	78
3.6	Experimental Analysis	79
3.6.1	Experimental analysis of the automated negotiation heuristics	80

3.6.2	Experimental results for differently sized graphs	82
3.6.3	Set-up and results for different graph structures	83
3.6.4	Experimental set-up and analysis of the collaborative filtering model	87
3.6.5	Results from retrieval experiments using cosine-based vs. correlation based similarity	89
3.6.6	Experimental results for selecting graph cut-off number of edges in the maximal graph	91
3.7	Discussion	94
3.7.1	Comparison to other automated negotiation (bargaining) approaches	94
3.7.2	Relation to graphical utility models and preference elicitation	97
3.8	Conclusions and future work	98

II Preferences under uncertainty and strategic reasoning in sequential auctions **103**

4	Designing bidding strategies in sequential auctions for risk averse agents	105
4.1	Introduction	105
4.1.1	Goals and organisation of this chapter	107
4.2	Modeling Utility Functions Under Risk	107
4.2.1	The importance of risk aversion in decision making: an example	110
4.3	Bidding in sequential auctions with complementarities	111
4.3.1	Optimal bidding policy for sequential 2nd price (Vickrey) auctions	113
4.3.2	Optimal bidding policy for sequential 1st price auctions: numerical solutions	115
4.3.3	Bidding strategy for multiple copy auction sequences	117
4.4	Experimental analysis	118
4.4.1	Experimental hypotheses	119
4.4.2	Experimental setup	119
4.4.3	Experimental results for one-type item auctions	120
4.4.4	Results for one item and different auction lengths	121
4.4.5	Setting with different item types and more complex preferences	123
4.4.6	Multiple item setting: hypotheses	125
4.4.7	Results for two-item case	127
4.5	Conclusions and further work	127
5	Using Priced Options to Solve the Exposure Problem in Sequential Auctions	129
5.1	Introduction	129
5.1.1	Options: basic definition	130
5.1.2	Related work	131
5.1.3	Outline and contribution of our approach	132
5.2	Expected profit for a sequence of n auctions and 1 synergy buyer	133
5.2.1	Profit with n unique goods without options	133
5.2.2	Profit with n unique goods with options	135
5.3	When options can benefit both synergy buyer and seller	137

5.3.1	When agents are better off with options	137
5.3.2	Synergy buyer's profit-maximizing bid	145
5.4	Simulation of a market with a single synergy buyer	151
5.4.1	Synergy buyer's bid strategy	152
5.4.2	Experimental results: market entry effect for one synergy buyer . . .	155
5.5	Settings with longer sequences of auctions and effect of auction order . . .	158
5.6	Multiple synergy buyers	160
5.6.1	Two synergy buyers interacting indirectly through the exercise price level	161
5.6.2	Direct synergy buyer competition in the same market	162
5.6.3	Larger simulation with random synergy buyers' market entry	163
5.7	Discussion and further work	164

III Emergence of collaboration and social preferences in web communities and online markets 167

6	Emergence of Consensus and Shared Vocabularies in Collaborative Tagging	169
6.1	Introduction	169
6.1.1	Tagging versus Taxonomies on the Web	170
6.1.2	Overview of Related Work	172
6.1.3	The Tripartite Structure of Tagging	175
6.1.4	Organization of the chapter	177
6.2	Detecting Power Laws in Tags	177
6.2.1	Power Law Distributions: Definition	177
6.2.2	Empirical Results for Power Law Regression for Individual Sites . .	178
6.2.3	Empirical Results for Power Law Regression Using Relative Frequencies	180
6.3	The Dynamics of Tag Distributions	181
6.3.1	Kullback-Leibler Divergence: Definition	181
6.3.2	Application to Tag Dynamics	182
6.3.3	Empirical Results for Tag Dynamics	182
6.3.4	Examining the dynamics of the entire tag distribution	184
6.4	Constructing Tag Correlation Graphs	185
6.4.1	Methodology	185
6.4.2	Constructing the tag correlation (folksonomy) graphs	186
6.5	Identifying tag vocabularies in folksonomies using community detection algorithms	189
6.5.1	Using community detection algorithms to partition tag graphs . . .	190
6.5.2	Edge filtering step	191
6.5.3	Normalized vs. non-normalized edge weights	191
6.5.4	The graph partitioning algorithm	192
6.5.5	Graph partitioning: experimental results	194
6.5.6	Eliminating tags from resulting partitions to improve modularity . .	195

6.6	Comparison benchmark: automatic construction of keyword vocabularies from search engine query data	196
6.6.1	Data set and methodology employed	198
6.6.2	Discussion of the results from the query log data and comparison	198
6.7	Conclusions and Future Work	200
6.8	Acknowledgments	201
7	The Complex Dynamics of Sponsored Search Markets: An Empirical Study	203
7.1	Introduction	203
7.1.1	The data set	204
7.2	Complex systems analysis applied to the web and economics	204
7.2.1	Power laws: definition	205
7.3	Influence of display rank on clicking behaviour	206
7.3.1	Results on display position bias and interpretation	207
7.4	Market structure at the advertiser level	209
7.4.1	Distribution of impressions vs. distribution of clicks for the top advertisers	209
7.4.2	Distribution of market share per display rank position	210
7.5	Using click data to derive search term recommendations	211
7.5.1	Deriving distances from co-occurrence in sponsored click logs	212
7.5.2	Constructing keyword correlation graphs	212
7.5.3	Graph correlation graphs: results	213
7.5.4	Automatic identification of sets of keywords	214
7.5.5	Discussion of graph partitioning results	215
7.6	Discussion	216
7.6.1	Contribution of the chapter & related work	216
7.6.2	Future work	217
7.6.3	Acknowledgements	217
IV	Conclusions	219
8	Conclusions and further work	221
8.1	Overview of the research contributions per chapter	221
8.2	Further work	224
8.3	Concluding remarks	227

V	An industrial application case	229
A	A Platform for Auction-Based Allocation of Loads in Transportation Logistics	231
A.1	Introduction	231
A.1.1	The multi-party logistics domain	232
A.1.2	Company profile	232
A.1.3	Automating multi-party logistics using agents	233
A.1.4	Goals of this work	233
A.2	Overview of the business case and our platform	235
A.2.1	Generating transportation orders	235
A.2.2	Computing prices and costs	237
A.3	Auction protocol and design of the auctioneer agent	237
A.3.1	Auction set-up	238
A.3.2	Auctions for loads with a short lead time	238
A.3.3	Auctions for orders with a longer time horizon	239
A.3.4	Total capacity of loads to be generated per day	239
A.3.5	Auctioneer user interface	240
A.4	Automated bidders: description and user interface	241
A.5	The carrier agents: description and user interfaces	242
A.5.1	Transportation model and carrier costs	243
A.5.2	Penalty for late deliveries	244
A.5.3	Information supplied about other carriers during the competition	244
A.5.4	Planning and bidding decision support interface	245
A.6	Outline of preliminary human bidding results	246
A.7	Discussion	247
	Bibliography	249

Chapter 1

Introduction

1.1 Agents and electronic markets

Multi-agent systems are one of the most promising new technologies to emerge in recent decades, at the crossroads between several fields such as artificial intelligence, distributed systems, economics and even sociology. Some authors [16, 231] have outlined a vision, in which many of the tasks performed today by humans are delegated to intelligent, autonomous and proactive programs, generically called software agents. A system composed of several such agents is called a multi-agent system (MAS).

Electronic markets represent key coordination mechanisms in multi-agent systems. They allow parties to efficiently allocate resources, tasks and capabilities in large distributed systems, composed of self-interested agents. The rapid rise in electronic commerce and marketing, logistics, distributed networks (among many others) have made the development of agent technologies capable of automating such processes increasingly important. For example, electronic commerce has witnessed an exponential increase in the value of the goods and services sold online just in the past few years. It is not just the sale of physical goods that has greatly increased, but also the sale of “virtual” services, such as screen attention space for displaying advertising in e-commerce, or keyword hits by surfers using search engines. Such sales require frequent, repeated interactions, which are the type of processes that are likely to benefit most from automation using software agents.

There are many challenges that designers of agents acting in electronic markets must face. Perhaps the most easily recognized challenge in designing and using such a system, is the lack of centralized control. Agents are autonomous actors, that take their own decisions, rather than simply executing operations assigned to them by an outside process (such as objects or web services do). Furthermore, perhaps more importantly, they are often self-interested actors, whose goals and objectives may not match. For example, in optimizing a transportation logistics network involving several carrier companies, the optimal alloca-

tion of loads from the perspective of each company may be very different than the optimal allocation for the entire system. In other application scenarios, such as online advertising, agents representing different companies actively compete for virtual commodities, such as consumer attention space.

An important challenge in MAS is the presence of uncertainty, i.e. incomplete or imperfect information, both regarding the market environment, the preferences, strategies and behaviour of the other agents and, sometimes, even uncertainty in specifying the agent's own preferences. Furthermore, unlike assumptions commonly made in game theory, the agents are bounded rational actors and often have to make decisions in limited time, under risk aversion or based on other constraints imposed by their owners or the market environment. Moreover, the opposing agents participating in the same market may also be bounded rational and even act 'irrationally', which makes modeling the agent's own "optimal" or "rational" behaviour in such a setting even harder.

Another important approach in the study of agent-mediated electronic markets are the so-called complex systems techniques. The aim of such approaches is to examine how order and structure can emerge in a large system composed of many autonomous entities (i.e. agents), acting independently, without any central controller to provide coordination. The recent surge of interest in systems such as web communities and online electronic markets, where structure emerges out of individual agent decisions, makes such questions increasingly important.

1.2 Negotiation (bargaining) vs. auction protocols

Negotiation, very broadly defined, is the "process by which a group of agents communicate with one another to try to reach agreement on some matter of common interest" c.f. [111, 189]. Automated negotiation has been at the forefront of research interests in the multi-agent research community ever since the beginning of the field [129, 189].

One of the main distinction lines being drawn in existing literature is between automated negotiation (bargaining) protocols and auction protocols [111, 147]. Bargaining is always a decentralized process and is typically (though not necessarily) based on an "alternating offers"-type of protocol. Some authors [111, 115, 175] (among others) argue that bargaining does have some advantages over auctions, especially in multiple issue cases, in which there is incomplete information about the opponent preferences (or even uncertainty about the agent's own preferences) and the space of possible deals to be explored is very large. Bargaining also allows more flexibility in how the negotiation is modeled, as well as a degree of self-interest on the part of the agents. Some sources [87, 115, 175, 179] even argue that, in electronic commerce, multi-issue negotiation should be modeled, at least partially, as a cooperative process, because sellers have an interest in maintaining a good relationship and the long-term satisfaction of their buyers.

Auctions, on the other hand, follow protocols with fixed rules, that typically rely on a trusted center to collect the bids and compute the winner (or winners) and corresponding

payments. Auctions have a long history of research in fields economics [38, 127, 158, 159], management science [190], but also artificial intelligence and theoretical computer science [27, 55, 58, 59, 81, 142, 174, 194]. They have been the method of choice for automating electronic marketplaces.

In our work, we have looked at both mechanisms, for different settings. Our initial work in the topic started on designing efficient bilateral negotiation mechanisms, first for linear utility functions (Chapter 2 of this thesis), then for complex, interdependent utility functions (Chapter 3). We have also looked at designing bidding strategies for sequential auctions, in settings not previously considered in existing literature, such as the case when some of the agents are risk averse (Chapter 4), or when options are auctioned instead of the items themselves (Chapter 5).

This thesis takes an engineering approach, meaning that we aim to identify open problems, and then engineer and validate solutions for them. We do study to what degree these problems can be addressed using an analytical, mathematical approach in so far as possible. However, many negotiation and auctioning processes are too complex to be solved using a purely analytical approach, as is normally the case for real-world problems. In such cases, experimental validation is a promising alternative, which was used extensively in this thesis.

A common thread running through the research presented in this thesis is that we take the heuristic approach to the design of bidding agents. That is, we focus our attention on designing the strategies that bidding agents use to bid or negotiate in a given market environment (usually one widely encountered in practice), not the market protocol itself. This is an important distinction, as explained in the next section.

1.3 Designing for strategic behaviour: market mechanisms vs. individual agent strategies

With the growth of interest in electronic markets, several research lines have emerged, proposing different approaches to modeling strategic, self-interested behaviour when allocating resources or tasks among a set of agents. One of the most promising such approaches is *computational mechanism design* - or to be more precise, that part of mechanism design theory that concerns design of electronic markets.

Mechanism design initially developed as a branch of algorithmic game theory [168]. Basically defined, mechanism design is concerned with defining the “rules of the game” (i.e. the market mechanism), such that the outcome (i.e. final allocation of the items, together with the corresponding payments) guarantees certain desiderata (i.e. properties). Commonly-cited desiderata include, for example: Pareto-optimality, efficiency, budget balance or individual rationality [58].

Besides from these game-theoretic desiderata, computational requirements (i.e. the computation time or memory needed to find such a mechanism) often play an important role. The mechanism design approach aims to relieve the need of strategic reasoning on the part of the

bidding agents, as the structure of the market provides bidders with an equilibrium bidding strategy. Different equilibrium concepts exist, varying in their strength, e.g. dominant strategies, ex-post efficient, Bayesian. The most desirable market mechanisms are strategy-proof mechanisms, i.e. those mechanisms in which truthful bidding is the dominant strategy.

The mechanism design approach has proven to be very successful in many applications. However, there exists a wide range of practical settings for which it is unrealistic to assume that one can design a completely new market mechanism from scratch. Furthermore, many mechanisms proposed by this line of research often involve allocation and bidding rules that are carefully designed and mathematically sound, but may be counter-intuitive for human users of the system.

Moreover, in many real-life allocation problems, there is more than one market an agent can/should participate in, and the strategic behaviour across “market borders” becomes the crucial issue. Even if an agent has an optimal (e.g. dominant) bidding strategy in each of the markets he participates in, when coordinating the bidding in different markets, his optimal strategy may be very different from the dominant strategy for each market taken in isolation. One such example is bidding in a sequence of second price (i.e. Vickrey) auctions, for agents which have complementary utilities over the items being offered [27, 89, 187, 217]¹. Although the agents in such a sequential auction have a dominant bidding strategy in each individual Vickrey auction taken in isolation, bidding optimally in a sequence of such auctions is a complex decision problem, and the bids placed in the optimal sequential bidding policy may differ considerably from the dominant bids in the individual auctions².

Similarly, a related decision problem is faced by agents bidding in a set of simultaneously ascending English auctions, when agents have complementary utility functions [1, 184]. Whereas a simple, dominant bidding strategy exists for each English auction taken in isolation, determining the optimal bidding strategy for the entire set is a challenging problem, for which no dominant strategy results are known. Intuitively explained, it is difficult for a bidding agent to distribute the additional complementarity value across a sequence or a set of simultaneous auctions, because an agent can only know if he can benefit from the complementarity once all items in the desired bundle have been acquired (i.e. once all the auctions close).

While the computational mechanism design community has begun to address some of these challenges, through such techniques as online mechanism design or adaptive mechanism design [33, 74, 92], these approaches still impose several restrictions on the structure of the problem, and for many market settings widely used in practice, no dominant equilibrium strategies are [yet] known to exist.

The research performed for this thesis mostly follows the other main direction of research

¹While these notions will be formally defined later, intuitively, a complementary valuation implies that an agent assigns a combination of items a super-additive utility (i.e. a utility higher than the sum of the items taken individually), while substitutability implies a sub-additive utility (see Chapter 3)

²In a sequence of auctions, this is true whenever either complementarity or substitutability effects exist between items, or there are other preference constraints to account for, such as budget constraints or aversion to risk (c.f. [27, 89, 187] and Chapters 4 and 5).

on agent-mediated electronic markets: given a market structure and protocols, how can one design the optimal agent strategies for bidding in such markets?

Work on designing bidding agents' strategies also has a relatively long history in the MAS community [1, 27, 81, 83, 89, 184, 188, 223] (among many others). Several platforms have been proposed to enable comparison of different auction trading strategies (as well as learning and adaptation heuristics). The most well-known is the Trading Agent Competition (TAC) platform, with its different versions: TAC classic, TAC supply-chain management etc. [96, 192, 227, 230]. The market structure for the TAC competitions is built to resemble trading scenarios that would be encountered in practice. Reasoning required for the trading agents in these platforms combines elements of both efficient bidding in sequential and simultaneous auctions, as well as learning, anticipation of future orders, inventory management etc. Another direction of work (mentioned here for completeness) examines bidding heuristics for double auction settings [226], which is characteristic to financial markets [124].

As already discussed, in this thesis we also take the heuristic approach, and we are mostly concerned about the design of agent strategies, rather than the market mechanism itself. In particular, we are concerned with one aspect of the problem, which is how to model and efficiently use *preference information* of the agents taking part in such markets.

1.4 Modeling preferences and utilities in agent-mediated market settings

In building efficient electronic markets, the method of modeling and reasoning about the preferences of participating agents is a key modeling choice. Some sources call modeling preferences of buyers and providers remains the “Achilles’ heel” in the application of multi-agent resource allocation to industrial procurement settings [46]. There are, however, considerable differences as to their meaning of the terms “preference” and “utility” in different sources in the economics and multi-agent system literature. In the broadest sense, preferences express the “relative or absolute satisfaction of an individual when faced with a choice between different alternatives” [46]. In this thesis, we broadly distinguish between two broad classes of concepts of preference or utility: preferences in combinatorial settings (i.e. to reason regarding multiple criteria or multiple items) and preferences under uncertainty.

Most existing literature on multi-agent resource allocation and market mechanisms considers combinatorial preferences and utilities. Combinatorial preferences are either multi-item (i.e. involve expressing preferences over combinations of items) or multi-issue or multi-attribute (e.g. involve combinations of attributes for the same item, e.g. colour, price and mileage for a used car, to take the example used in Chapter 2). Moreover, in many (if not most) realistic applications, it is reasonable to expect that there are complex dependencies between attributes or items, and the choice in one may affect the choice made for a subset of others (i.e. the preferences are non-linear). Efficient representation and reasoning with such

non-linear, combinatorial preferences in market situations is a complex problem, which we discuss in Sect. 1.7 below.

Another perspective on defining preferences considers the complex decisions an agent faces, not so much with respect to specifying desired combinations over multiple items or issues, but with respect to uncertainty about the future. This appears to be a standard understanding of “preference” in some fields, such as econometrics. For instance, in a 2006 MIT textbook on econometric analysis of auction data [171], the chapter regarding preferences deals exclusively with preferences towards risk. In this thesis, we consider both perspectives on preference, and both types of market-based interaction discussed above: bilateral negotiation and auctions, although for different problems.

1.5 Emergence of collaboration and structure in multi-agent systems

As noted in Section 1.1, another important problem arising in multi-agent systems is the lack of centralized control. Nevertheless, many systems occurring in real-life that one would intuitively recognize as “multi-agent”, exhibit a remarkable degree of structure, although they lack any recognizable central authority or “controller”. Instead, order seems to emerge from the decentralized actions of many autonomous agents, acting independently to satisfy their own interest. Examples of such systems include: the formation of equilibria and pricing structure in markets (a phenomena first referred to by Adam’s Smith as the “guiding hand”), emergence of stable vocabularies in human languages (but also in tagging systems) [40, 93, 206], formation of stable groups in online social networks [131] etc. This raises questions not only regarding the existence and properties of such stable structures, but also the dynamics of the process, i.e. how do they form, especially in an environment with no central information source and/or self-interested parties.

One of the recently emerging fields that aims to study such phenomena is complex systems theory [11, 160, 228]. The seminal work of Robert Axelrod on the evolution of cooperation [7] marked a turning point, since it showed, through computer simulations, how cooperation can emerge in a multi-agent system, even in the absence of a central authority.

A related discipline that aims to examine complex-systems type phenomena through large-scale simulations is agent-based computational economics [218]. There has been much work recently in this area. For example, researchers have simulated the dynamics of artificial agent societies [7], stock markets [6], and even entire economies [37, 68]. The development of the web has given a new stimulus to this work, and researchers have built complex simulations of the emergence of social networks [131], online market systems [54] or artificial languages and semiotic dynamics [42, 210].

But, perhaps more important, the emergence of the “social web” provides, for the first time, the opportunity to test these hypotheses empirically, on real-world data. In fact, while there has been a lot of work on simulations (starting from different assumptions), only re-

cently did researchers begin to study this problem at an empirical level, using large scale datasets generated by the actions of very many (thousands, or in some cases even millions) of web users.

Interestingly, many of the effects found resembled closely what was hypothesised in complex systems theory from the beginning. In particular, it appears that there are important “network effects” when many users collaborate online and make decisions in an online community or marketplace. What this means, basically, is that the actions and choices made by previous users may considerably influence the choices made by future users. This type of self-reinforcing feedback loop often gives rise to the so-called “power law” distributions [11, 41], which are characteristic of large-scale systems that can be characterised as “complex”.

This thesis makes two important contributions to understanding the emergence of social structure in such large-scale, decentralized systems. One is collaborative tagging (results presented in Chapter 6) and sponsored search markets (Chapter 7).

1.6 Positioning of the contributions of this thesis

The previous discussion identified some important open challenges in understanding and designing multi-agent systems:

- Complexity of representing (and reasoning with) complex preferences. These include both combinatorial preferences and preferences towards risk and uncertainty.
- Strategic reasoning of agents based on these complex preferences, especially for cases when agents are self-interested.
- Lack of central control, and especially, the emergence of cooperation in the absence of a central authority.

In this thesis, we aim to make several contributions to the state of the art in understanding, modeling and solving these challenges, as follows.

Part I of the thesis is mostly concerned with the issue of modeling combinatorial preferences (multi-issue or multi-item) in bilateral negotiations. Chapter 2 considers how preference information can be efficiently used in a negotiation model in which users preferences are expressed over several discrete attributes and one continuous attribute (price). Chapter 3 considers how complex, multi-issue negotiations over many binary items or bundles of items can be modeled using utility graphs. Part I also deals with some issues related to strategic reasoning, since, although bilateral negotiation is often a partially cooperative process, there is an important degree of self-interest involved on the part of the bargaining agents.

Part II of the thesis can be seen as mostly dealing with preferences in uncertain environments and strategic reasoning, in particular the strategic reasoning of agents with inter-dependent valuations participating in a sequence of auctions. Chapter 5 discusses how a

priced options mechanism can help solve the exposure problem bidders face in a sequential auctions setting (and implicitly, the strategic reasoning it involves during the bidding process). But the issue of preference is also important in this part as well, although in the form of agent preferences towards risk, when faced with an uncertain future. In this context, Chapter 4 examines how an agent's preferences towards risk affects his/her optimal bidding policy and resulting market allocation.

Finally, Part III of the thesis studies the issue of lack of centralised control, more particular the emergence of collaboration and structure in a large multi-agent system, in the absence of a central controller. For this part, we use large scale, empirical data from two important “social web” applications: collaborative tagging and sponsored search markets. While in collaborative tagging, for instance, the issue of strategic/game-theoretic reasoning does not play a direct role (since there is no competing allocation of some scarce resource), still the issue of how agents take decisions is a crucial one to model. Arguably, there is also a connection to the issue of preference, since through their choice of tags and links to click agents express an implicit opinion (which may or may not be influenced by that of other users).

The rest of this introductory chapter is organised as follows. In the following sections (Sect. 1.7-1.10), we give more detailed descriptions of the problems in this space which we aim to address in this thesis, as well as brief abstracts of our results for each. In Section 1.11 we give the overview of the structure of [the rest of] the thesis. Section 1.11 also summarizes the structure of the thesis through a diagram, such as to more intuitively highlight and explain the relations that exist between the different chapters. The introduction concludes with a list of resulting refereed publications related to each chapter.

1.7 Modeling of combinatorial preferences (multi-issue or multi-item) in bilateral negotiations

There are many ways to express a choice between multiple outcomes defined in multi-agent, economics and AI literature. A taxonomy of preferences used in the multi-agent literature would include:

- *Qualitative preferences*: No numerical utility values are assigned to outcomes, only value labels such as “good”, “very good”, “unsatisfactory” etc.
- *Quantitative preferences*: Preference over outcomes are expressed in the form of a utility function (to be defined below). Note that sometimes qualitative and quantitative preferences taken together are called “cardinal” preferences.
- *Ordinal preferences*: only an order can be specified between ranking (i.e. through an asymmetric and transitive preference relation between alternatives).
- *Fuzzy*: - a degree of preference can be specified for each alternative, etc.

In this thesis, we generally consider quantitative preferences, i.e. those preferences for which a user can assign a numerical value (either utility or monetary value) to possible outcomes (or combination of items), and the discussion in the following sections of this introduction refers to the quantitative case.

However, note that, if one defines preference as any choice between several outcomes or alternatives, the concept of preference can be constructed as broader. For example, actions such as choosing a tag that other users have also selected in the past (see Chapter 6), or clicking on the link at the top of a list, in order to save reading time (see Chapter 7) may be seen as expressing an implicit preference. We return to this idea in Section 1.10, when describing the chapters in Part III of the thesis - the discussion in the following sections referring to the case of quantitative, economic preferences.

The basis of quantitative preference modeling is utility theory. Following the work of Kenney and Raiffa [179], many multi-issue negotiation and resource allocation models use a utility function, which maps the outcome space over a set of issues (attributes, criteria) to a utility value, which is frequently - though not necessarily - scaled between 0 and 1. The crucial thing to note is that Raiffa's models and much of the initial research on multi-issue negotiation considers *linearly additive utility functions*, i.e. each issue/attribute under negotiation is assigned a weight, and the utility of each possible outcome/contract is computed as a weighted sum over the issues under negotiation.

A special subclass of quantitative preference functions, which is implicitly used in most of existing auction literature, are the so-called "quasi-linear" preferences. This basically means that the utility of the agents is expressed in monetary terms (as an amount of money), as opposed to utils (i.e. conventional units, usually scaled between 0 and 1). This can be viewed as a restriction for some settings, as real utility functions over monetary endowments are known to be concave, i.e. humans are known to have a decreasing marginal utility for money (see, e.g. [171]).

1.7.1 Pareto-optimal outcomes in multi-issue negotiation

As shown in [179], multi-issue and multi-attribute negotiation models are fundamentally different from single-issue negotiation (such as bargaining over a price). Multi-issue negotiations represent non-zero sum games, in the sense that it is possible to find mutually beneficial trade-offs between the issues under negotiation such as to increase the gains for both parties. Raiffa also shows that the more asymmetric preferences between the negotiators are, the higher the potential for mutually beneficial trade-offs between the issues.

The main criteria to measure how efficient an agreement (or contract) is the so called Pareto efficiency. An outcome is said to be Pareto optimal if it is not strictly dominated by any other outcome in the preferences of *both* (or all) sides (agents) in a negotiation. That means, there are no trade-offs possible that would increase the utility of one agent, without making another agent worse off. The set of all Pareto-optimal points form the so called Pareto-frontier. The distance of an outcome to this frontier gives a measure of how efficient

that outcome/contract is.

An important concept in designing multi-issue negotiation models is the uncertainty regarding opponent preferences, defined here as the amount of information regarding the opponent preferences available when making negotiation offers.

Direct vs. indirect revelation mechanisms

The literature on agent-mediated electronic markets identifies two main approaches by which agents can share their private preference information:

- *Direct revelation mechanisms.* Direct revelation mechanisms are based on the revelation principle [58]. Basically explained, the revelation principle states that any allocation mechanism with a certain equilibrium can be transformed into another mechanism, in which a trusted center asks the agents to truthfully reveal their preferences and implements the original equilibrium (and allocation) on their behalf. This means that typically (though not exclusively), direct revelation leads to a centralized allocation mechanism, such as a combinatorial auction.
- *Indirect revelation mechanisms:* In this type of mechanism, the agents are not assumed to directly reveal their preferences to the other agents, but communicate their preferences throughout their counter-offers (or their bids, for an auction). For instance, in a bilateral, multi-issue negotiation over the sale of a car (see [115, 116] and Chapter 2 of this thesis), the agents do not directly reveal to each other how much they are willing to pay to get their favourite colour or their favourite accessories (e.g. CD player, air conditioning) installed, but in practice, this can be deduced indirectly from the offers/counter-offers they make. Similarly, an agent representing a customer on a large electronic commerce website (see [185, 186, 220] and Chapter 3) does not have to reveal all his preferences to the merchant, but the merchant (who may or may not also act as the auctioneer) can learn part of their preferences from previous counter-offers.

Basically, in this thesis we take the indirect revelation approach, as we argue this is more realistic in many real-life applications, in which only a limited degree of trust exists between parties in sharing information and no fully trusted third party can be established. The reasons for this may be endogenous to the negotiation mechanism (e.g. there is no “optimal” incentive compatible mechanism and the opponent may use any information supplied to get a better deal for himself) or exogenous (e.g. it may be undesirable to have to specify preferences over the whole set of alternatives, due to privacy concerns or future business interests).

Furthermore, for complex non-linear preferences, there are also preference formulation and communication costs. As we show in Chapter 3, because of bounded rationality and communication ability, it is often difficult for an agent herself to formulate and communicate bids over all possible item combinations in advance.

In the remainder of this introduction, we will discuss some important directions of research into modeling preferences in agent-mediated markets, followed by brief descriptions of the contributions made in this thesis to open problems in the field. Therefore, some sections of the introduction present important concepts from a general point of view, while other sections describe how these general concepts were extended by our own research, described in the chapters of this thesis. The goal is to allow the reader to get a better understanding of the positioning and contribution of our work with respect to the state of the art in the field.

1.7.2 Modeling multi-attribute negotiation with incomplete preference information

An important direction of work in the literature on multi-issue negotiation is how to design efficient bargaining strategies in settings when agents do not have any information about the opponent's (i.e. negotiation partner's) preferences. They may have, however, some prior knowledge about the domain they are negotiating about. This prior domain knowledge can be, for example, fuzzy logic distances between attributes, such as the perceptual distance between different colours (such as in [71, 163]), or an ordering between a set of qualitative attribute labels (such as "good", "standard", "meager" etc), in our research (see Chapter 2).

The work presented in Chapter 2 and [115, 116] considers such an incomplete information negotiation model. The aim of this model is to investigate the role that partially revealing preference information can improve the outcome of a multi-attribute negotiation. As a practical domain case, we considered a bilateral negotiation between a buyer (customer) and a seller (car dealer) over the sale of a car. The negotiation is not exclusively on price, but also on the quality of the accessories which the dealer has to install in the car to get the deal done (such as a CD player, extra speakers, air conditioning and tow hedge). In this setting, we show that it is possible for both parties to reach close to Pareto-efficient agreements, by revealing only partial (i.e. incomplete) information about their preferences of the negotiation partner. Furthermore, we proposed a novel guessing heuristic, by which an agent uses the history of opponent's bids to predict his/her preferences in order to propose better deals.

1.7.3 Non-linear and combinatorial preferences in negotiation

A crucial problem in applying multi-issue or multi-item negotiation models in many realistic settings is the fact that there may be complex inter-dependencies between different issues, leading to non-linear preferences or utility functions [108, 126, 138, 186]. The problem appears both when considering integrative, multi-issue negotiations, as well as negotiations over bundles of items [46, 186]. In both these cases, it is important to allow for concise representations of the utilities over possible outcomes³.

³The problem is in fact, two-fold. First there is the complexity related to *preference formulation* of combinatorial preferences, as well as one of *preference communication* complexity.

The easiest way to represent preferences is to enumerate all possible outcomes (or combinations of goods), together with their utility value for those goods (monetary or otherwise). This is called the *explicit form* of preference representation (or “bundle form”). The explicit form is fully expressive, in the sense that any utility function may be described by listing all possible combinations and their values. It is, however, impractical for most non-trivial settings, as the number of descriptions would be exponential in the number of resources (e.g. for only 50 binary issues or items, $2^{50} > 10^{15}$ values would need to be assigned - see Chapter 3). This has prompted another important direction of research in electronic markets, that of designing more concise *utility representation (or preference) languages*. There are several classes of such preference languages:

- Bidding languages, which are typically used in combinatorial auctions to allow agents to formulate their bids (and, implicitly, communicate their preferences to the auctioneer). Some specific bidding languages include:
 - The OR language: The agent can specify an array of valuations over different subsets of items in a given bundle of items. The value of any combination can then be computed as the maximal value that can be obtained as a sum over disjoint subsets specified [46]. For example, in the bid: $\langle \{I_1\}, 3 \rangle OR \langle \{I_2\}, 3 \rangle OR \langle \{I_3\}, 3 \rangle OR \langle \{I_1, I_2\}, 8 \rangle$ expresses that the bidder is willing to pay 3 for either I_1, I_2, I_3 or 11 for all 3 items (in this case, it is better to take the value of the subset $\langle \{I_1, I_2\}, 8 \rangle$ than the values of each individual item separately). Because the OR dependency is not exclusive, the OR language cannot express substitutability dependencies, i.e. it cannot express the fact that getting a combination of items has lesser utility than the sum of individual items. In the above example, it is not possible for the agent to express that he is only willing to pay 4 if he gets both I_1, I_2 . If the bid $\langle \{I_1, I_2\}, 4 \rangle$ were added to the set of bids placed, then the auctioneer would simply match the bids over the individual items $\langle \{I_1\}, 3 \rangle$ and $\langle \{I_2\}, 3 \rangle$ (as any terms of the OR dependency may be chosen).
 - “Exclusive OR” (i.e. XOR) bidding language [194] - is an alternative to OR, in which all combination bids are assumed to be mutually exclusive. For example, in the above example, a bid such as: $\langle \{I_1\}, 3 \rangle XOR \langle \{I_2\}, 3 \rangle XOR \langle \{I_3\}, 3 \rangle XOR \langle \{I_1, I_2\}, 4 \rangle$ means that the agent (bidder) can either use only one item from I_1, I_2, I_3 with a utility of 3, or the combination of $\{I_1, I_2\}$ with a utility of 4, but no other combination (so, e.g. getting both I_1 and I_3 would still only have the utility of 3). XOR is fully expressive, in the sense that it can represent any monotonic utility function. However, XOR may have a high communication/elicitation cost, even for simple settings. An example is the utility function that, for any set R of items $I_1, \dots, I_n \in R$, simply counts the number of items the agent owns - i.e. $u(R) = |R|$. Such a function would require an exponential number of bids in the XOR language, but only a linear number in OR language. This is because, using XOR, all combinations specified are mutually exclusive, so all possible combinations of the n items of size $|R|$ needs to be

specified as a terms in the bid. Because of this issue, and because OR language is considered more natural way to represent preferences, there exists a line of work that aims to extend its expressiveness, without requiring an exhaustive listing of XOR bids [169].

- Weighted propositional formulas and straight line programs are other alternatives to representing complex preferences, which make use of logical formalisms. For example, weighted propositional formulas are derived from a qualitative form of preference representation, in which the preferences of the agent are expressed as goals. In the weighted case (unlike in purely formal logic approaches), goals can be assigned a utility weight if satisfied. We do not deal with this kind of preference languages in this paper, but the reader can consult [139] for the full details of this approach.
- The k -additive form [46, 55, 186] (also called the polynomial form [133]) is another natural and concise method to represent combinatorial preferences. K -additive functions can encode synergy (complementarity or substitutability) effects between subsets of up to k items. For instance, if we denote by x_1, \dots, x_n the instantiation of the set of n items, the expression for a 3-additive utility form (i.e. taking a maximum $k=3$) is:

$$U(x_1, \dots, x_n) = \sum_{1 \leq i \leq n} \alpha_i x_i + \sum_{1 \leq i, j \leq n} \alpha_{i,j} x_i x_j + \sum_{1 \leq i, j, k \leq n} \alpha_{i,j,k} x_i x_j x_k \quad (1.1)$$

Where x_1, \dots, x_n represents a vector of 1 and 0, denoting whether an item is (or is not) considered in the combination being evaluated, the reals $\alpha_1 \dots \alpha_{n,n,n,n}$ are the parameters of the function, while k (same k as in "k-additivity") is the maximum rank of the polynomial, i.e. all the polynomial terms having a rank above k_{max} have the coefficients $\alpha = 0$. Linearly additive functions form a subclass of the k -additive class, as defined above, for $k_{max} = 1$. The k -additive form is fully expressive, for unbounded k . This means that, if k is sufficiently large, it can be used to express any utility function over a given, finite, binary set of items. In practice, although (as discussed in Chapter 3 this thesis), in order for this representation to be computationally useful, the maximum rank of the polynomial k is generally assumed to be bounded to a limited value (e.g. 2-4, as discussed in Chapter 3).

1.7.4 Modeling multi-item negotiations over k -additive utility functions using utility graphs

In Chapter 3 of the thesis, we consider the case of modeling complex bilateral negotiations over a set of multiple, binary issues (which can also represent a bundle of items). From the concise representation forms discussed in the taxonomy from Sect. 1.7.3, the one we used, as we found it most natural in the context of the multi-issue negotiation, is the k -additive form. This representation is a natural extension of linear utility models already used in much previous multi-issue bargaining literature although, as we will show, allowing any degree on

nonlinearity in preference makes the bargaining problem considerably harder. For example, in Eq. 1.1 above, the case of $k_{max} = 2$ is already much harder than $k_{max} = 1$.

Multi-issue negotiation with non-linear utility functions is known to be a complex problem, even for the case of binary issues [108, 126, 207]. The state of the art in this field proposes complex solutions that involve a mediator, as well as techniques such as simulated annealing [126] or econometric methods [207] that are either computationally expensive or do not scale well for settings with many issues. In [186] (corresponding to the first part in Chapter 3) we introduce a novel utility graph formalism for modeling nonlinear (i.e. k-additive) preferences, and we show how such graphs can be used to model and learn opponent preferences in complex, multi-issue negotiations. Utility graphs are originally inspired from probabilistic graphical models, but they encode utilities, rather than probabilities. The main idea behind our approach is to use the structure of the graphs to restrict the opponent modeling and search to the most promising region of the utility space. A seller agent can start a negotiation with an approximation of the utility function of a typical random buyer in the form of a maximal utility graph, and then refine this model based on the counter-offers he observes during the negotiation. In our case, the initial utility graph reflects the prior information that the seller has about how the utility function of a random buyer is structured, in order to help in the search.

An important question is, of course, how does the seller acquire this initial buyer utility graph approximation. One solution is to assume some prior domain knowledge, such as plausible constraints on the shape a utility function could take (which may be reasonable for some settings). For e-commerce domains, we have proposed another alternative: using collaborative filtering on previous sales data, that will be presented after the discussion in the next subsection.

1.7.5 Individual preferences and social influence

In the previous discussion on combinatorial preferences, preferences are defined from a single-agent perspective, meaning that the utility of any agent is assumed to be private and *independent of what other agents may desire*. Otherwise put, if a seller encounters a buyer and negotiates with him the configuration of a product or the composition of a bundle of items, he will assume that the preferred combinations of this particular buyer are completely independent of what other buyers encountered wanted in the past. This is, in fact, a standard assumption in much of negotiation and auction theory.

However, existing practice in electronic commerce suggested, for the Chapter 3 of this thesis, an alternative approach. The success of social search in providing online buying recommendations provides considerable evidence that preferences are not strictly independent, but are in some way clustered. Consider for example, the case of Amazon.com, who has several million book titles in its collection. Eliciting, for each individual customer, his/her preferences over these books to propose acceptable bundles for the buyers would be a nearly impossible task. However, Amazon implicitly assumes that if a large numbers of customers that bought a certain book in the past were also interested in another book (or set of books),

then there is a high probability that future customers may be interested in this combination as well. For instance, if a customer buys a book on travelling to Portugal, the Amazon engine assumes he may also be interested in a book on travel to Spain, since many customers encountered in the past showed interested in both. Therefore proposing a deal (e.g. postage reduction, or a small discount) may be a good way to incentivise the customer to buy both books from the site.

Note that this does not always have to be a correct prediction: in fact many customers may not be interested in the exact combination proposed. However, it does provide a good approximation in searching the space of a customer's preferences, even if the customer was never encountered before.

Traditionally, research in multi-issue negotiation does not explicitly model this social dimension of customer preferences, or consider the role that social influence plays on the structure of utility functions. We show that having an explicit representation that relates the two fields (in our case in the form of utility graphs) allows us to considerably improve search in an online negotiation setting. Furthermore, the interaction between these fields does not have to be one-way: negotiation also has a lot to add to web-based recommendation in electronic commerce. Through an iterative negotiation process, the initial proposals (based on anonymous, aggregate preferences) can be customized to the preferences of a particular customer, based on the indirect revelation made through his/her counter-offers in negotiation.

1.7.6 Learning the structure of utility graphs used in multi-item negotiation through collaborative filtering

Our approach to modeling opponent preferences in negotiation makes use of the above intuition. Chapter 3 of this thesis proposes a novel collaborative filtering method by which previously concluded negotiation data can be used to construct the initial approximation of the utility graph of a random buyer that the seller can use in later negotiations. The seller will then adjust (learn) the values in the graph, for each specific negotiation, based on the counter-offers the buyer makes, until an agreement is reached over the bundle combination. Therefore, we take what can be described as a two-step customization approach: initially, an approximation of the maximal structure of a utility graph for a random buyer is obtained using collaborative filtering on all concluded negotiation data (which does not have to be buyer-specific). Then, this deal is refined through offers and counter offers during the negotiation with a specific customer.

We show that the combined approach can enable buyers and sellers to reach efficient agreements even in complex non-linear settings, involving only indirect revelation (although there are some assumptions regarding the maximal complexity of the utility graphs that a buyer can have). One of the contributions of this approach to the state of the art in automated negotiation is that it provides a link between the customization techniques used in multi-issue or multi-item negotiation and those used in collaborative filtering and social computing. In fact, we show that the use of collaborative filtering techniques can lead to more efficient

and shorter negotiations for complex, non-linear utility settings than was reported in other research [126].

1.8 Preferences under uncertainty and bidding in sequential auctions

In the previous discussion, we have mainly discussed the concept of preference (or utility) in the context of integrative negotiation, in which the allocations for all items (or issues) is agreed at the same time. Thus, when an agent specifies a preference by assigning a monetary value to a combination of items, he is bidding for an entire combination, and there is no uncertainty that he will not get some of items in the agreed configuration, if the seller accepts the offer. This is a reasonable assumption for integrative negotiation and combinatorial auctions (where the allocation for all items is negotiated simultaneously). However, it does not hold for other widely used allocation mechanisms, such as sequential/simultaneously ascending auctions [27, 89, 184, 217] or one-by-one issue negotiations [72].

In this section (corresponding to Chapters 4 and 5 of the thesis), we consider the case when agents have to bid sequentially items sold in different auctions, without knowing with certainty that they will get the entire combination of items they desire. In such cases, economic theory identifies another important class of preferences, preferences towards risk. Risk aversion is a very important part of economic theory - in fact, a 2006 MIT textbook on the econometrics of auction data [171], the chapter on “preferences” is basically concerned with preferences towards risk.

The way econometric theory models risk aversion is through the so-called Neumann-Morgenstern preference functions, in which the utility derived by an agent from a certain amount of money is not a linear function, but a concave one. Otherwise stated, utility functions are not quasi-linear, in the sense that the utility that each agent derives from an amount of money is not directly proportional to the amount paid/received.

In the following, we briefly define the exposure problem in sequential auctions, the role that risk aversion plays in the bidding decision, as well as an overview of the contributions in Chapters 4, 5 and Appendix A.

1.8.1 Sequential auctions and the exposure problem

As shown in Sect. 1.3 above, there are two main directions of research in the application of agent systems to auction markets. One concerns the design of the auction mechanism itself, such that participant agents have a dominant bidding strategy (usually, to declare truthfully their values), as well as certain properties, such as efficiency, individual rationality or budget balance. However, for many market designs that are necessarily encountered in practice, such as sequential or simultaneously ascending auctions, this is not possible, and research

has focused instead on designing the bidding strategies of agents participating in such auctions.

As shown in [27, 89, 187] and Chapters 4 and 5 of this thesis, the main problem that a bidder faces in a sequential auction is the exposure problem. Informally defined, the exposure problem means that an agent has to commit to buying an item, before he/she can be sure that he will be able to secure other items in his useful set or bundle (defined as the set of items that gives him a positive utility). If she fails to acquire this bundle, then he makes a loss. Hence, we say that the agent is exposed to the risk of a loss.

Most of the models that study bidding auction bidding start from the assumption that agents have quasi-linear utility functions. Basically defined, quasi-linearity assumes each agent has a set of payoffs that he/she assigns to any combinations of items. These payoffs are, for many of the models studied, private: they are not known to the other parties. The utility that an agent gets from participating in the auction is assumed proportional to the difference between his/her private payoff and the amount he pays to acquire the items in question, in other words, it is defined strictly in monetary profit/loss terms.

This quasi-linearity of preferences assumption, while widely used and valid for many business models and settings, does not universally hold. In many real-life settings, even assuming it is true that agents have private values for different subsets of items under negotiation, profit and loss are not judged in the same terms. Making a loss from an interaction (i.e. paying more than his/her private payoff value) is not proportional as gaining the same amount as profit. In other words, agents are risk-averse to making a loss, even if the potential for gain is considerably larger.

1.8.2 Designing sequential auction strategies for risk-averse bidders

Prior to the publication of our research, there had been quite a lot of previous work on designing efficient bidding strategies for agents participating in sequential [27, 89, 217] and simultaneously ascending [1, 184] auctions.

While this work reported some positive results, an important limitation of existing literature we examined was that it does not explicitly model the risk-taking attitude of the bidding agents. By “explicitly model” we mean building a profile of the agent’s risk preferences towards uncertain, future outcomes (such as the final allocation of a sequential auction). In standard economic theory, since the seminal work of K. Arrow and J. Pratt, preferences towards risk have been considered essential in understanding and modeling decision making under uncertainty [5, 88, 153, 171]. Auction literature from standard economics [158, 171] considers risk aversion an important problem in modeling real bidder preferences. However, the economic literature that we are aware of does not consider sequential auctions with complementary bidder valuations, except perhaps in the simplest of settings (because such auctions do not have well-defined equilibria). More specifically, unlike the AI community, researchers in economics are not concerned with designing automated bidding heuristics for sequential auctions.

The main contribution of Chapter 4 of this thesis is making a link between risk-aversion models, and the strategies that risk-averse bidders can use in sequential auctions. First, we introduce the Arrow-Pratt risk models from economics to the problem of modeling agent bidding strategies. We then study the way in which the perceived optimal bidding strategy computed by a risk averse agent, given her probabilistic model of the future, differs from the optimal strategy of a risk neutral agent. We find that agents more averse to risk bid more aggressively, in order to cover their sunk costs for the initial items in the sequence. However, if the future sequence of auctions is initially perceived as too risky (given the agent's initial estimation of future closing prices), the best strategy available to a risk averse agent is simply not to participate at all.

Our experimental results show that a risk averse bidder has, as expected, a lower chance to end up with an incomplete bundle of goods, thus make a loss. However, when considering long-term and repeated interactions, such agents make, on average, a lower expected profit, because they participate in less auctions. For some market settings, this also affects, in a negative way, the auctioneer revenues from the auctions.

In the following section, we look at a different side of the problem of exposure to risk of loss in sequential auctions, namely what can be done to reduce it.

1.8.3 Options mechanisms in sequential options

As discussed above, sequential auctions do not guarantee a dominant bidding strategy for the agents (unlike the combinatorial case). However, the problem remains, as many allocation problems occurring in practice are inherently decentralized and sequential. Different sellers may prefer, for a variety of reasons, to sell their items separately - or even through different markets, as the number of electronic auction sites online indicates. Furthermore, in many application settings, not all resources that are to be allocated are known in advance, but they appear dynamically over time. In Chapter 5 of this thesis, we study an alternative to this very difficult problem that, although it cannot completely eliminate bidder's exposure, it can significantly reduce it: the use of priced options.

Intuitively defined, an option is a contract between the buyer and the seller of a item, where the buyer has the right to choose in the future whether or not he will purchase the item against the pre-agreed *exercise price*. The seller is then bound to sell the item at the demand of the buyer. Since the buyer gains a right, he has to pay the *option price* regardless of whether he will exercise the option or not.

Options reduce the exposure problem a synergy buyer faces. He still has to pay the option price, but if he fails to complete his desired bundle, then he does not pay the exercise price as well, and thus he limits his loss. The risk of not winning subsequent auctions is partly transferred to the seller, who may miss out on the exercise price. However, the seller can benefit indirectly from the participation in this market of additional complementary-value buyers (also called "synergy buyers"), who would have otherwise stayed out.

Our work builds on an idea first proposed in Juda & Parkes [120, 121]. Their work pro-

poses a market design with free (i.e. zero-priced) options. To prevent buyers from hoarding options (i.e. which they may have an incentive to do, as they are offered free and can be exercised optionally), bidders are assumed to place their bids only through proxy agents provided by the mechanism. They show that, in this market mechanism, truth-telling is a dominant strategy on the part of the buyers. The sellers are incentivised to use the proposed options mechanism by market entry effects. We note that there are some important limitations to their approach, which we aimed to address through our work. First, market entry effects are often not sufficient to motivate the sellers to offer options for free (due to the risk of remaining with their items unsold) and, in such cases, only positively-priced options can provide sufficient incentive for both sides to use the mechanism. Moreover, the design proposed in [120] cannot deal with the case when several synergy buyers are active in the market simultaneously. With priced options, while the problem of setting exercise prices becomes more difficult with multiple synergy buyers present simultaneously in the market, options can still be shown to be beneficial.

1.8.4 Using priced options to solve the exposure problem

By comparison to [120, 121], although it starts from the same idea, the work described in Chapter 5 takes a different approach. Rather than attempting to design a complex, custom-made mechanism, our goal is to investigate, in a decision-theoretic model, under which conditions selling options for the items would be more beneficial for both sides in a market (sellers and buyers), by comparison to direct auctions.

We consider a model in which buyers obtain the right to buy the item for a certain exercise price in the future. Each option is described by a fixed exercise price and a flexible, market-determined option price. The seller fixes the exercise price of an option for the item he has for sale and then sells this option through a first-price auction. Buyers bid for the right to buy this option, i.e., they bid on the option price. Note that in this model, direct auctions appear as the particular case of fixing the exercise price at zero: such options would always be exercised, assuming free disposal.

Our approach and analysis can be characterized as decision-theoretic, meaning both buyer and seller reason with respect to expected future prices. First, we consider a setting in which n complementary-valued items are auctioned sequentially, assuming there is only one synergy buyer (the competition consists of local bidders desiring only one item). If options are auctioned for these items instead of the items themselves, the agent may bid an option price corresponding to a higher total amount (option + exercise price) than in a direct sale, because he does not have to pay the exercise prices if he fails to get the desired combination. However, the seller also takes an exposure to the risk of a possible loss by auctioning options instead of items, because the buyer may not exercise the acquired options, and hence he/she may not collect the exercise price. Therefore, in order for him to have an incentive to offer options, he expects an increase in the bids he receives to compensate for this risk. For this setting, we show analytically that using priced options can increase the expected profit for both the synergy buyer and the seller, compared to auctioning the items directly.

Second, we study experimentally market settings in which multiple synergy buyers are active simultaneously. In such settings, the problem of fixing the right exercise price becomes harder, because the seller has to maximize expected buyer participation, but at the same time reduce his own exposure. We find that, while some synergy buyers loose because of the additional competition, others may actually benefit, because sellers have an incentive to fix exercise prices at high enough levels which encourage more synergy buyers to participate in the bidding.

1.9 Applications to transportation logistics

Besides the theoretical results discussed in Chapters 2 - 5, we also investigated a practical business setting in which sequential auctions are used: distributed transportation logistics. As this research represents a practical case study, rather than a contribution to fundamental research, the full details are included as Appendix A, rather than a chapter. However, the experience gained from performing this case study served to inform many of the choices made in our theoretical work on auction bidding strategies.

Transportation logistics and supply chain management represents a challenging, but potentially very fruitful area for the application of agent-based electronic market techniques, such as the ones we consider. The increasing complexity and shifting structure of modern supply chains, as well as increasing competitive pressures in this market has led to an increasing demand and interest for such distributed optimization techniques, involving multiple parties. The practical impact of improved allocation which can be achieved through such techniques can be significant. For example, in the Netherlands, the average transport performance is about 60%-70%. Improving this utilization rate is also the goal of the DEAL (Distributed Engine for Advanced Logistics) project, which groups together several universities and large logistics service providers in the Netherlands. The applied research work reported in Appendix A was carried out in collaboration with Vos Logistics Organizing, Nijmegen, one of the largest European transportation logistics companies.

1.9.1 Auction-based allocation of transportation loads in multi-party transportation logistics

Multi-party transportation logistics settings are those in which the company that accepts an order to transport goods does not necessarily also own the actual capacity (i.e. trucks) to also carry out this order. This is an increasing trend in modern supply chains, where often multinational companies with large, regular amounts of cargo to be delivered prefer to outsource these orders to other companies that undertake to find convenient delivery options, within a set of pre-negotiated terms. These intermediary logistic companies then negotiate how to distribute these orders with other, often smaller, companies who have the actual transportation capacity (which own the actual trucks and hire the drivers). This is actually a cheaper and more efficient option in many cases, as a fit for the order can be searched

in the transportation plans of several companies. In standard transportation management literature [221] such distributed supply chains are called multi-party logistics.

Typically, such supply chains are composed of companies with different business models. Third-party logistic providers (3PL companies) are those that have their own transport capacity (i.e. truck fleet) and plan this own capacity, while fourth party logistic providers (4PL companies) are those which “orchestrate” the supply chain, i.e. acquire large sets of orders from large shippers and then re-distribute these orders among a set of other companies with actual transport capacity. This is also the way in which Vos Logistics Organizing operates. Our goal in the joint project was to examine how such outsourcing activities could potentially be automated, by building an auction-based platform, in which transportation orders are allocated among a set of agents representing different companies.

Thus, the business case that this platform addresses involves a large 4PL company receives orders from a large shipper (orders can arrive throughout the day) and has to find transportation capacity for them among a set of other 3PL companies. The 4PL company that outsources the orders acts as the auctioneer, while the 3PL companies are the bidders taking part in a set of (reverse) auctions, in which the bidder that offers the lowest transportation costs usually wins (although there is some flexibility and the exact auction design in the decision of the auctioneer, see Chapter A). Our initial goal was not to design specific bidding strategies, but to build an auction platform around a real transportation business case, in which actual bidding strategies (both human and automated) can be implemented and tested. As a next step to this practical user study, the strategies developed for repeated allocation in our more theoretical work (such as risk-averse bidding strategies, option mechanisms etc.) could be re-implemented in this platform, and also tested against the bidding strategies used by human planners. The full details of the implementation, choices made and results are presented in Appendix A.

1.10 Preferences in social web communities and online markets

The last part of this thesis, consisting of two chapters, is about modeling collective preferences in large web-based systems. These chapters are different in their methodology and scope from the others, in the sense they are mostly concerned with empirical analysis of large-scale web data, rather than proposing theoretical models and validating them through simulations. Nevertheless, as we show in Chapter 6, this research is highly relevant to the issue of modeling preferences in online settings, including electronic markets. Although the concept of “preference” is used here more generically, as a choice between several alternatives/outcomes, rather than the narrower, utilitarian sense taken in economics, the topic is relevant to the topic of the thesis. Moreover, it often uses similar techniques, such as the graphical models. This part also complements the approach taken in other chapters, by providing analyses of empirical web data produced from the actions of thousands (or even millions) of human users, rather than computer simulation data.

The techniques we use for this analysis are so-called “complex systems”-type techniques. Basically stated, the goal of such techniques is to examine how the local actions of individual agents can lead to the emergence of order or structure at the global level [11]. For online environments, the user action usually consists of a click on some link or ad, or assigning a tag to a resource, while the emerging structures can be correlation graphs between terms, based on the clicks they receive, market power of the advertisers in the system etc. Complex systems analysis usually allows insights into phenomena such as feedback and the influence of previous actions on other users.

This Section is divided into two separate topics: analysis of the dynamics of tagging systems, and an analysis of sponsored search markets.

1.10.1 The complex dynamics of collaborative tagging systems

The first part concerns a large scale analysis of collaborative tagging data. This line of research resulted from a collaboration after an extended stay at the Santa Fe Complex Systems Institute. It resulted in several papers, including a WWW’07 paper [93] and an extended journal version.

In this joint work, we use data from the social bookmarking site *del.icio.us* to empirically examine the dynamics of tagging systems, and especially whether coherent categorization schemes can emerge from the unsupervised tagging done by many individual users. Our findings are two-fold. First, we find that final tag frequencies describing most tagged resources converge to “power law” distributions. This is important, as it shows that coherent, stable categorization can actually emerge in such a system, without a central controller of the terms which should be used. The specific power law-type distribution has been shown to appear in other complex, decentralized systems, and is typically interpreted to show that an implicit feedback process is at work in the system [166]. Other than studying the shape of the final distributions, We also propose an information-theoretic method to examine the dynamics of this convergence to stable distributions, both for the most used tags, and for the so-called “long tail” of tag distributions.

The second part of the chapter proposes a method to construct information structures from collaborative tagging, which we call tag correlation (or “folksonomy”) graphs. In such graphs, the distances between nodes (representing individual tags) is indirectly proportion to their co-occurrence frequencies in describing the same resources. It is worth noticing that the methods used to derive these tag correlations graphs closely resemble the collaborative filtering methods used for the utility graphs developed in our negotiation work (see Chapter 3 of the thesis).

We also describe a method to partition such graphs to obtain simple tag vocabularies, using so-called “community detection” algorithms. In this context, the simple vocabularies that can be obtained from applying such techniques on tagging data are compared to those that can be extracted from search engine data for the same domain. We find that the methods applied to collaborative tagging data outperform the results from a large search engine, at

least in some respects.

1.10.2 An empirical analysis of sponsored search markets

A final direction of research, also using empirical techniques on large scale web data, is reported in Chapter 7, which examines the dynamics of sponsored search markets. Sponsored search, in which advertisers pay for the clicks received by their text ads displayed alongside search engine results, has become an important part of modern advertising. It now represents the main source of revenue for large search engines, such as Google, Yahoo or Microsoft. Furthermore, since advertising link slots are allocated through an auction, which takes place automatically for each placed query, sponsored search provides fertile ground for investigating automated market-based techniques in a real, large-scale setting.

The data used for the study in this Chapter comes from the search engine Live.com⁴. Basically, this study, like the study in Chapter 6 also takes an complex-systems type of approach to analyzing sponsored search data. In this context, we show that not only the search keywords themselves, but also the relative weights of the advertisers in the market follow power law distributions. Thus, a small number of advertisers have a dominating share of the number of user clicks in the market. We discuss how this effect may be due to the way winners are determined, based on their historical click-through rates, which reinforces the position of the top advertisers through a feedback-type mechanism. Furthermore, we also study the issue of user attention, especially how users prefer to scan a list of sponsored search results. We find a clear relationship between the position occupied by a text ad in the list returned to a user and its probability of receiving a click.

Finally, we use our graph-based and collaborative filtering techniques (similar to the techniques as in Chapter 6, but now based on advertising click data), to output recommendations on which sets of search terms are most commercially promising. Here, in computing similarities between search terms, we only consider the queries that lead to an actual click on a returned text ad. Note that this is different from using organic search or tagging results, since we only take into account the “opinion” of a subset of users that show a clear interest in buying something online.

1.11 Structure of the thesis

Following the outline described above, this thesis is organized in 6 chapters, divided into three main parts. Part I contains two chapters that both deal with modeling multi-issue or multi-item preferences in complex, bilateral negotiations. Chapter 2 proposes an agent architecture that can handle incomplete preference information in negotiations over several discrete-valued attributes and one continuous attribute (price). Chapter 3 presents a negotiation model that uses utility graphs to represent many binary-valued issues, but with interde-

⁴This data was kindly provided to us by Microsoft Research, as part of a “Beyond Search” grant.

pendent, k -additive utility functions. Furthermore, Chapter 3 also discusses how the initial structure of a buyer's utility graph can be learned from past data, based on collaborative filtering techniques.

The second part of the thesis consists of two chapters that deal with auction bidding strategies, especially bidding strategies in sequential auctions. Chapter 4 shows how an agent's attitude towards future risk influences his/her optimal decision-theoretic bidding strategy, when participating in multiple, sequential auctions. In Chapter 5, we consider a different approach to solving the difficult problem of sequential allocation, and we show how priced options could reduce the exposure to risk that bidders with complementary valuations face in such auctions. Appendix A describes an auction platform that would allow bidding strategies (both automated and human) to be tested in a practical scenario: allocation of orders in transportation logistics.

The third part of the thesis deals with emergent collaboration and large-scale empirical studies of web systems. Chapter 6 studies how stable distributions and information structures can emerge in decentralized tagging systems, while Chapter 7 uses similar, complex-systems type approaches for an empirical study of sponsored search markets. The thesis is concluded by a discussion of the results and a discussion of issues that were left to further work.

Figure 1.1 summarizes, in graphical form, the basic structure of the thesis, and it highlights (through arrows) the connections that exist between different parts and chapters. These connections are either at the level of the fundamental problem being studied, or at the level of the technique used to address those problems. For example, both Parts I and II consider different types of agent-mediated market techniques (negotiations or auctions), especially problems related to modeling complex preferences and strategic reasoning in such settings. The two chapters in Part I both deal with the bilateral negotiation problem, while the two chapters in Part II deal with the problem of sequential allocation in competitive environments, in particular sequential auctions.

There are also important connections between chapters that may seem initially unrelated in scope, at the level of the type of techniques used. For example both Chapter 3 and Chapter 6 use similar graphical models techniques (in fact ideas developed in the initial research for Chapter 3 served at the source of inspiration for the tagging work from Chapter 6). These two chapters also share important idea about how social preferences can be exploited and how collaborative filtering can be used to learn graphical structures from past data. The two chapter in Part III are directly related by their use of complex systems type techniques, applied to empirical web data, although for different settings (tagging systems vs. sponsored search markets).

But there are also connections at the level of practical applications of the techniques in different chapters. For example, the applied case study work reported in Appendix A actually served to motivate and guide the type of design choices and problems to be studied in our more fundamental work on auction bidding, described in Chapters 4 and 5. One of the reasons we focused on the sequential auction case in these chapters is that sequential auctions appeared very frequently in practical case studies we examined, yet we found few theoretical results for this setting in prior literature.

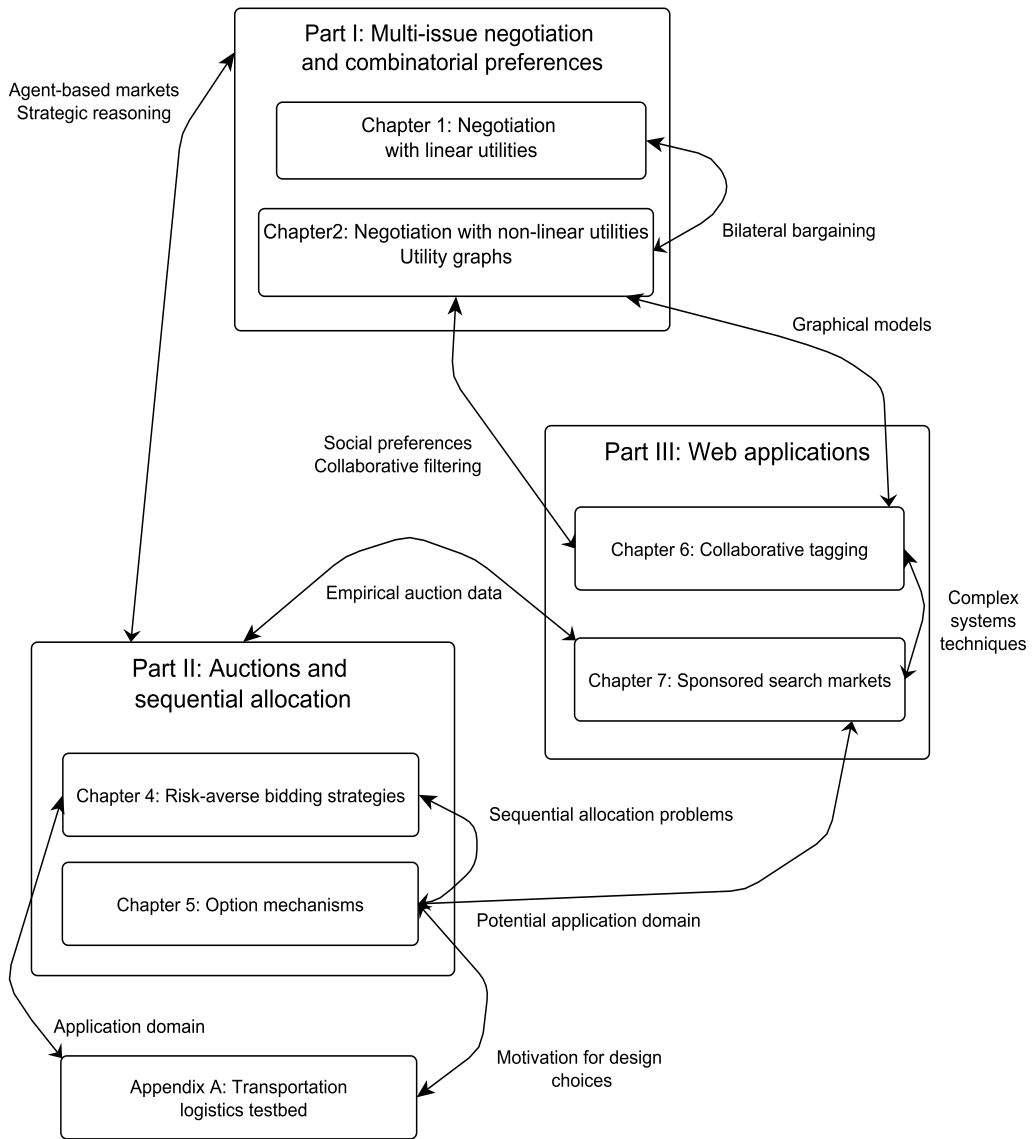


Figure 1.1: Thesis structure and relations between the work presented in different chapters of the thesis

Finally, there is also a connection between the chapters in Part II and the keyword auc-

tions in sponsored search markets. The keyword auctions from sponsored search provide large scale, empirical data for the type of sequential auction bidding behaviour, studied at a more theoretical level in Chapters 4 and 5. Chapter 7 provides a first indication of the type of insights that could be extracted from such data. Furthermore, we envisage that some techniques, such as the use of options, could potentially also be used for sponsored search markets, as repeated / sequential bidding problems are also important in those settings.

For the reader, we note that all chapters of this thesis, while they do refer to each other at different points, are self contained and can be read independently. In the remainder of this introduction, we provide a list of the most important publications that resulted from each of the chapters, as well as brief references to other published research results of the author that were not included in this thesis.

1.12 Publications related to each chapter

All the Chapters of this thesis are based on at least 2 peer reviewed publications (journal, refereed international conference or refereed book chapter), as follows:

- **Chapter 2:** Jonker, C., Robu, V., Treur, J. “An Agent Architecture for Multi-Attribute Negotiation Using Incomplete Preference Information”, *Journal of Autonomous Agents and Multi-Agent Systems (JAAMAS)*, Springer-Verlag, Vol. 15(2), pp. 221-252, Oct. 2007. A conference version of this paper appeared as: Jonker, C.M., and Robu, V., “Automated Multi-Attribute Negotiation with Efficient Use of Incomplete Preference Information”, *Proceedings of the Third International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'04)*, New York. ACM Press, 2004, pp. 1056-1064. A book chapter describing the platform used for both this work and other negotiation work performed at VU Amsterdam appeared as: Bosse, T., Jonker, C.M., Meij, L. van der, Robu, V., and Treur, J. “A System for Analysis of Multi-Issue Negotiation”, In: *Calisti, M., Klusch, M., and Unland, R. (eds.), Software Agent-Based Applications, Platforms and Development Kits*, Birkhauser Publishing, Springer-Verlag Group, 2005, pp. 253-280.
- **Chapter 3:** Robu, V., Somefun, D.J.A., La Poutré J.A. “Modeling Complex Multi-Issue Negotiations Using Utility Graphs”, *Proc. of the Fourth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'05)*, Utrecht, The Netherlands. ACM Press, 2005, pp. 280-287. The collaborative filtering extension first appeared as: Robu, V., La Poutré J.A. “Retrieving the Structure of Utility Graphs for Multi-Issue Negotiation Through Collaborative Filtering of Aggregate Buyer Preferences”, *Proc. of Second International Workshop on Rational, Robust and Secure Negotiations in Multi-Agent Systems (Best Paper Award)*, Springer Series on Computational Intelligence, vol. 89, 2008, pp. 147-168. An earlier version was presented in 2005 at the *Proc. of the Pacific Rim Workshop on Multi-Agent Systems (PRIMA'05)*, but the Springer LNCS post-proceedings were only confirmed to appear in 2009. A journal version, comprising both parts is under submission.

- **Chapter 4:** Robu, V., La Poutré, J.A - “Designing bidding strategies in sequential auctions for risk averse agents: a theoretical and experimental investigation” *Post-proceedings of the 9th Int. Workshop on Agent Mediated Electronic Commerce (AMEC'07)*, Honolulu, Hawaii, 2007. Springer Lect. Notes in Business Informatics, vol. 13, pp. 76-89. An extension of the results appeared as: Robu, V., La Poutré, J.A - “Designing risk-averse strategies for bidding for outsourcing through sequential auctions” *Advances in Agent-based Complex, Automated Negotiations - Proc. of the 1st Int. Workshop (ACAN'08)*, Estoril, Portugal, Springer Series on Computational Intelligence (to appear, 2008).
- **Chapter 5:** Mous, L., Robu, V., La Poutré, J.A - “Using Options to Solve the Sequential problem in Sequential Auctions”, *Post-proceedings of the 10th Workshop on Agent Mediated Electronic Commerce (AMEC'08)*, Estoril, Portugal. Springer Lecture Notes in Business Informatics (LNBI), Springer-Verlag (to appear, 2009). A synopsis of the main results was also reported as: Mous, L., Robu, V., La Poutré, J.A - “Can Option Mechanisms Solve the Exposure Problem in Sequential Auctions?”, *ACM SIGEcom Exchanges*, vol. 7, no. 2, ACM Press, July 2008. A more extended version, which includes the detailed proofs and full result graphs for different aspects (and which is the version included in this Chapter) appears in the *Proceedings of the Dagstuhl Seminar on Multi-Agent Planning Systems*, DROPS Electronic publication, Nov. 2008.
- **Chapter 6:** Halpin, H., Robu, V., Shepherd, H.- “The Complex Dynamics of Collaborative Tagging”, *Proceedings of the 16th International World Wide Web Conference (WWW'07)*, Banff, Canada 2007, ACM Press, pp. 211-220. The extended, journal version of this paper has been conditionally accepted as: Robu, V., Halpin, H., Shepherd, H. - “Emergence of Consensus and Structured Vocabularies in Collaborative Tagging Systems”, *ACM Transactions on the Web*, ACM Press (conditionally accepted, 2009). A preliminary version also appeared as: Halpin, H., Robu, V., Shepherd, H.- “On the Dynamics and Semantics of Collaborative Tagging”, *Proc. of the 1st Semantic Authoring and Annotation Workshop (SAAW'06)*, Atlanta, Georgia, USA, 2006 and as a technical report of the *Santa Fe Complex Systems Institute*, Santa Fe, New Mexico, USA.
- **Chapter 7:** Robu, V., La Poutré H., Bohte, S. - “Analyzing data from sponsored search markets from an agent-based and complex systems perspective”, *International Workshop on Agents and Data Mining Interaction (ADMI'09)*, Springer Lecture Notes in Artificial Intelligence (16 pages, to appear 2009-2010). A journal special issue version is in preparation.
- **Appendix A:** Robu, V., Noot, H., La Poutré, J.A., van Schijndel, W.J. “An Interactive Platform for Auction-Based Allocation of Loads in Transportation Logistics”, *Proc. Of 7th Int. Joint Conf. on Aut. Agents and Multi-Agent Systems (AAMAS'08)*, Industry Track, Estoril, 2008, IFAAMAS Press, 2008, pp. 3-11. A preliminary version describing the underlying business case appeared as: Putten, S.v.d., Robu, V., La Poutré, J.A.,

Jorritsma, A., Gal, M.- “Automating Supply Chain Negotiations using Autonomous Agents: a Case Study in Transportation Logistics”, *Proc. Of 5th Int. Joint Conf. on Aut. Agents and Multi-Agent Systems (AAMAS'06)*, Hakodate, Japan, 2006, ACM Press, 2006, pp. 1506-1513. The demonstration platform was also nominated as a finalist for the *Amsterdam Science Park Innovation Prize*.

Furthermore, papers resulting from almost the chapters were presented at the Belgium-Netherlands Conference on Artificial Intelligence (BNAIC) as extended abstracts.

1.12.1 Research work not included in the current thesis

Besides the work reported above, the research work of the author also lead to a few other publications that have not been included here. For completeness, we provide references to it and brief abstracts here, but the reader can skip this part, without loss of understanding of the other parts.

Decommitment mechanisms in transportation logistics

- 't Hoen, P.J, Redekar, G., Robu, V., La Poutré, J.A. “Decommitment in a Competitive Multi-Agent Transportation Setting”, In: *Calisti, M., Klusch, M., and Unland, R. (eds.), Software Agent-Based Applications, Platforms and Development Kits, Whittestein Series on Software Agents Technologies*, Birkhauser Publishing, Springer-Verlag Group, 2005, pp. 409-433. A shorter version appeared as: 't Hoen, P.J, Redekar, G., Robu, V., La Poutré, J.A. “Simulation and Visualization of a Market-Based Model for Logistics Management in Transportation”, *Proc. of the 3rd Int. Joint Conference on Aut. Agents and Multi-Agent Systems (AAMAS'04)*. ACM Press, 2004, pp. 1218-1219.

Decommitment is the action of foregoing of a contract for another (superior) offer - and has been extensively studied, for the bilateral case, by [4, 193, 195]. In the above book chapter [216], we study the decommitment concept for the novel setting of a large-scale logistics setting with multiple, competing companies. We consider a setting in which transportation orders are allocated to agents representing trucks through a system of distributed auctions. Using a large scale, spatial simulation, this work shows that, if agents are allowed to decommit from acquired transportation orders as better opportunities appear, the mechanism can lead to significant increases in expected profits. One-sided decommitment method is related to the option method reported in Chapter 5 of this thesis. The difference is that in the model described in [216], trucks are only allowed to pass orders to each other through decommitment if they belong to the same company. This removes strategic reasoning problems associated with two-sided decommitment, which makes the problem different in nature. The book chapter cited above also describes a novel visualization platform was developed to illustrate the decommitment strategy, for different input scenarios.

Design of a PLAT trading strategy

- Silaghi, G.C., Robu, V. “An Agent Strategy for Automated Stock Market Trading Combining Price and Order Book Information”, *Proc. of the ICSC Congress on Computational Intelligence Methods and Applications (CIMA 2005), Advanced Computing in Financial Markets Track*, Istanbul, Turkey, IEEE Press, 2006.

This work was performed in collaboration with G. Silaghi on designing a mixed automated stock trading strategy (i.e. one that combines several strategies) for the Penn-Lehman Automated Trading Platform (PLAT) [124]. Our strategy combined several existing strategies, based on a weighted “vote” function.

Workshop proceedings edited

- Ito, T., Zhang, M., Robu, V., Matsuo, T. (editors) “Rational Robust and Secure Negotiations In Multi-Agent Systems”, *Proc. of IAT/WI Workshops*, IEEE Press, 2008.
- Ito, T., Zhang, M., Robu, V., Fatima, S., Matsuo, T. (editors) “Advances in Agent-Based Complex, Automated Negotiations”, volume in *Springer Series on Computational Intelligence*, Springer-Verlag, 2009.

Part I

Modeling combinatorial preferences in bilateral negotiation (bargaining)

Chapter 2

An Agent Architecture for Cooperative Multi-Attribute Negotiation With Incomplete Preference Information

2.1 Introduction

Recent years have seen a surge of interest in negotiation technologies, seen as a key coordination mechanism for the interaction of providers and consumers in future electronic markets that transcend the selling of uniform goods [134, 211]. Suggested applications range from modeling interactions between customers and merchants in retail electronic commerce [91], to the online sale of information goods [208], or reducing operational procurement costs of large companies [14]. Such technologies could prove especially useful in the case of multi-attribute negotiations, which represent non-zero sum games, where “as values shift along multiple directions, it is possible for both parties to be better off” [189].

Gutman and Maes [91] discuss the difference between competitive and co-operative negotiation models in electronic commerce. They show that modeling retail market negotiations as strictly competitive assumes that merchants are unnecessarily hostile to customers and, furthermore, it offers them no long-term benefits. They conclude that sellers often care less about profit on any given transaction and care more about long-term profitability, which implies customer satisfaction and long-term customer relationships. Their analysis makes a strong case for co-operative negotiation models for the retail market.

However, even assuming partial cooperation, in many application settings such as e-commerce, only a limited degree of trust exists between parties in sharing preference in-

formation. The reasons for this may be endogenous to the negotiation (e.g., fear the other may abuse this information to get a better deal) or exogenous (e.g., privacy concerns). This position is supported, among others, by [71] who argue that “what is required are agent architectures that implement different search mechanisms, capable of exploring the set of possible outcomes under both limited information and computation assumptions”.

The work presented in this chapter starts from the basic intuitions described above: that in multi-attribute negotiations in electronic commerce often agents are only willing to share partial preference information to each other directly (or even none, for some cases). Therefore, agents have to use any amount of preference information the opponent is willing to share, and if not, to build a model of opponent preferences based on his/her bids or counter-offers.

Our work considerably extends an agent architecture and negotiation model developed in the AI department of the Free University in Amsterdam [118], to allow it to deal with incomplete or partial information. We take into account two different types of incomplete information:

- Partial profile information, which is communicated by the negotiation partner herself in the beginning of the negotiation.
- Profile information which can be deduced (learned) from successive bids during the negotiation itself. Here we start from the assumption that the way the negotiation partner is bidding may reveal something about his preferences. For this mechanism we use the term *guessing* to clearly show it is a heuristic.

The Chapter is organised as follows. First, Section 2.2 presents the formal design of the negotiation model. Next, the experimental validation of the model is presented in Section 2.3. Section 2.4 discusses the results of our model and compares it other approaches in literature, while Sect. 2.5 concludes the chapter with a discussion.

2.2 The multi-attribute negotiation model

The bilateral negotiation considered in our model follows an alternating-offers protocol. A bid in such a negotiation has the form of values assigned to a number of attributes.

To make it more intuitive, the multi-attribute negotiation model in this chapter is built around a specific domain example: the sale of a car. In this domain, the relevant attributes to be considered are: CD player, extra speakers, airco, tow hedge, price. A bid then consists of an indication of which CD player is meant, which extra speakers, airco and tow hedge, and what the price of the offer is. Although the model and examples are built around this domain, the underlying negotiation techniques proposed are generic, and this section provides a generic formal description of the model. Instantiations in other domains are possible

and have been considered - for example an employer and employee negotiating about work shifts and overtime pay (work performed in collaboration with Almende B.V., Rotterdam).

As already mentioned in the introduction, the model in this chapter builds on results initially presented in [118]. This initial model was adapted, through the work of the main author of this thesis, in two main directions (c.f. [115, 116]):

- A mechanism where the agents are allowed to exchange and take into account partial preference information from the negotiation partner was modeled.
- A novel heuristic by which an agent can estimate the preferences of the other using his past bids was proposed and tested (we call this a “guessing” heuristic).

Both for the original work and the extension, the DESIRE software environment and conceptual design methodology [31] were used to design the agents. DESIRE was a long running research project at VU Amsterdam, in which conceptual designs of multi-agent systems for many application areas were developed. It is beyond the topic of this thesis to go into the specifics of the DESIRE methodology, especially as it was not used later in other chapters. The reader should note, however, that some decisions taken in this chapter, for example the separation of the algorithm into conceptual components, were initially motivated by the use of the DESIRE method.

We do cover, however, some elements of the original, DESIRE-based negotiation model, to allow more extensive explanations for the parts that were added or adapted through the research of the author. For further details readers are asked to consult [118] and [116].

The proposed negotiation model works by performing computations on two levels: the overall bid level and the attribute level. This involves first evaluating the utility opponent’s previous bid, and then planning the target utility for the own next bid. Finally, the configuration of the next bid will be selected such that it fits this target value. In the design of our agent, these steps are modeled as separate DESIRE components (see Fig. 2.1), and our presentation follows this structure.

2.2.1 Bid Utility Determination and Planning Component

The evaluation for each attribute is computed based on an evaluation function, specified by the agent owner (user) in the beginning of the negotiation. This function takes the generic form $\text{eval}: V \rightarrow E$, where V is either a finite set of discrete values or continuous interval, while $E = [0, 1]$. For example, in the car sale domain, accessories have discrete values (quality levels, assigned an evaluation by the user), while attributes such as mileage or price are continuous.

The utility of the negotiation partner’s previous bid is computed by first computing the utility of each attribute, and the overall bid utility. The overall utility U_B for each bid or contract combination B is taken as a weighted sum of the attribute evaluation values E_j for the different attributes (issues) j :

$$U_B = \sum_j w_j E_j \quad (2.1)$$

From the above formula, one can already see that the model in this chapter assumes *linearly additive utility functions* between the attributes. Here, all weights w_j are normalized importance factors based on the raw importance factors ϑ_j for the different attributes. The user (owner) of each agent provides these non-normalised importance factors ϑ_j through an interface, in the beginning of the negotiation. These are then normalized as:

$$w_j = \frac{\vartheta_j}{\sum_k \vartheta_k} \quad (2.2)$$

For each consequent bid an agent makes, first a target evaluation is chosen at the overall bid level. For determination of the next bid's target utility TU the following formula is used: $TU = U_{BS} + CS$, with U_{BS} is the utility of the agent's own last bid ("bid of self"), and the concession step CS determined as:

$$CS = \beta \left(1 - \frac{\mu}{U_{BS}}\right) * (U_{BO} - U_{BS}) \quad (2.3)$$

Where U_{BO} denotes the utility of the opponent's last bid ("bid of other"), but computed with respect to the agent's own utility function. Note that this choice for the concession step follows the existing negotiation model [118] and was designed to ensure that the target utility for the next bid is always scaled between $[0..1]$. Factor β stands for negotiation speed, while factor $(1 - \mu/U_{BS})$ expresses that the concession step will decrease to 0 if U_{BS} approximates a minimal utility μ . The minimal utility μ is a measure of how far the agent is willing to concede to the opponent in the current negotiation round.

2.2.2 The Attribute Planning Component

This component (whose internal decomposition is shown in Figure 2.1) determines the attribute values for the next bid, in such a way that the next bid will always have the target utility as its utility. This is done in two steps: first a target evaluation is computed per attribute, based on the target evaluation planned for the whole bid. Next, attribute values are chosen with the evaluation closest to the target evaluations, for all attributes except price. The configuration of the next bid is then completed by selecting a value for price, such that the utility of the final bid fits exactly its target. In order to make better directed concessions, in planning the target evaluation for each attribute we take into account not only the own preference weight of the agent, but also the weight of the opponent. If the opponent is not willing to reveal her preference weight for some (or maybe all) attributes, an estimation of these weights is computed in "Estimation of Opponent's Parameters" component. The role of the "Guess Coefficients" component is to analyze the way the opponent is bidding and to provide some extra information to be used for estimating these private preference weights. In the following we discuss these components in a separate sections.

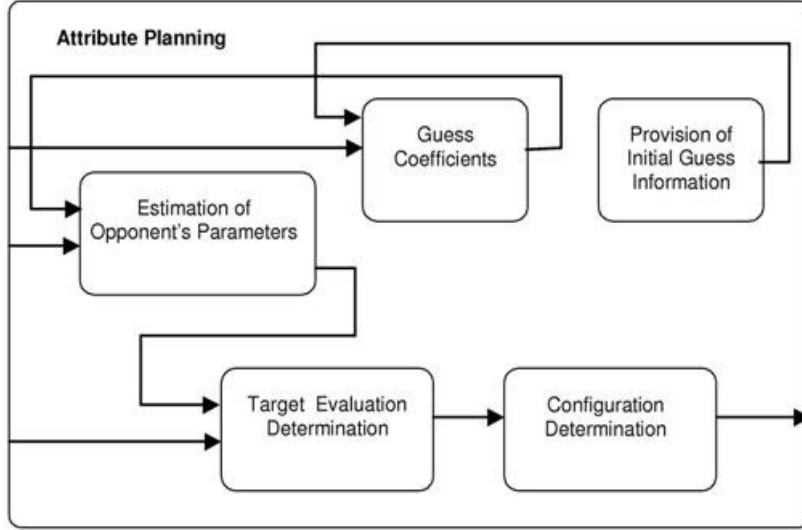


Figure 2.1: Internal composition of Attribute Planning

2.2.3 The Target Evaluation Planning Component

This component outputs a target evaluation for each attribute in the next bid, based on the bid target value. The target attribute evaluation is determined in two steps. First a basic target attribute evaluation for each attribute is computed as:

$$BTE_j = E_{BS,j} + \frac{\alpha_j}{N}(TU - U_{BS}) \quad (2.4)$$

Where $E_{BS,j}$ represents the evaluation for attribute j in the agent's own previous bid, U_{BS} the overall evaluation of the agent's previous bid (i.e. "bid of self"), while TU represents the target utility for the next bid (computed as shown in Section 2.2.1). The parameter α_j is chosen as $\alpha_j = (1 - w_j)(1 - E_{BS,j})$, where the first parameter expresses the influence of the user's own importance factor, while the second factor assures that the target evaluation values remain scaled in the interval between 0 and 1. Parameter N is a normalization factor, defined as: $N = \sum_j w_j * \alpha_j$. By this choice we ensure that the following relation holds: $\sum_j w_j BTE_j = TU$ (for a full proof of this property we refer the reader to the paper describing the original negotiation model [118]).

The Basic Target Evaluation, however only takes into accounts the own preference weights of the agent. Using only this value does lead to a working model, but the outcomes that it leads to can be sub-optimal, as the preferences of the negotiation partner (or opponent agent) are not taken into account when computing concessions. To improve on this, the following solution was implemented.

Based on this existing negotiation framework, we can now propose a model to deal with the incomplete information. For each attribute $j \in A$ (where A denotes the set of all attributes) a *Preference Difference Coefficient* δ_j is computed as:

$$\delta_j = \frac{W_{other,j} - W_{own,j}}{W_{other,j} + W_{own,j}} \quad (2.5)$$

This coefficient (scaled between -1 and 1) expresses how different the preferences of the two parties for each attribute are. Positive values for δ_j denote a stronger preference of the negotiation partner for attribute j , while negative values denote a stronger own preference for this attribute.

The concession to be made in each attribute $j \in A$ depends on a parameter called *configuration tolerance*, denoted as $\tau_j \in [-1, 1]$. The tolerance parameter is chosen to be attribute-specific, in order to better differentiate the amount of concessions between attributes. Therefore, for each attribute $j \in A$, the configuration tolerance depends on the preference difference coefficient of that attribute, according to the following formula:

$$\tau_j = \tau_{gen} * (1 + \delta_j) \quad (2.6)$$

Where the parameter τ_{gen} represents the general tolerance, used by the agent for all attributes j . The general tolerance is always chosen between 0 and 0.5 and also gives a measure of how fast the agent is willing to make concessions. Values closer to 0 will denote an agent who is less willing to make concessions, while values closer to 0.5 will denote an agent who is interested to reach a deal quickly. Since $\delta_j \in [-1, 1]$, the tolerance for any attribute j is scaled between 0 and $2 * \tau_{gen}$.

Finally, the target evaluation for each attribute j is computed. This is done by taking into account both the basic target attribute evaluation (as described above) and a concession to the attribute evaluation from the previous bid of negotiation partner, as follows:

$$TE_j = (1 - \tau_j)BTE_j + \tau_j E_{BO,j} \quad (2.7)$$

Where BTE_j is the basic attribute evaluation for attribute j and $E_{BO,j}$ is the evaluation for attribute j from the opponent's previous bid. From the above formula, one can see that values of the configuration tolerance τ_j close to 0 signify that mostly the user's own importance factors are taken into account, while values close to 1 shows that maximum possible concession is made towards the other's value. And since τ_j depends directly on δ_j , it is the difference in preference for each attribute that determines how much concession is made in that attribute.

Because, in this model both the sum of the agent's own weights and sum of the opponent's weights are always scaled to 1, the above mechanism leads to a situation where greater concessions in some attributes (more important to the opponent) will always be balanced by smaller concessions in other attributes (more important to him/herself). Such an

asymmetric concession system allows both negotiating parties to reach agreements closer to Pareto-optimality, because each party gets a higher local utility value for the attributes he/she cares more about.

In this component we have assumed that the opponent's preference weights for an attribute are known. However, if the other is not willing to share his weights for some (or all) attributes, then they will need to be estimated.

2.2.4 Estimation of Opponent's Parameters Component

This component determines, for those attributes for which the opponent was not willing to reveal his preference weights, an estimation of those weights.

First, we denote by A_{known} the set of attributes for which the opponent was willing to reveal his importance weights in the beginning of the negotiation and by $A_{unknown}$ the attributes whose preference weights are kept private. Since all preference weights are normalised (c.f. Section 2.2.1), the sum of weights for the private attributes is computed as:

$$\sum_{j \in A_{unknown}} W_j = 1 - \sum_{k \in A_{known}} W_k \quad (2.8)$$

For attributes with private weights, the remaining weight $\sum_{j \in A_{unknown}} W_j$ has to be divided between them. For this purpose we assign a parameter called the *Remaining Weight Distribution Coefficient* $R_j \in \mathbb{R}$ to each attribute $j \in A_{unknown}$.

The attributes of unknown weight can be further classified into two subsets:

- Attributes for which a reliable guess about the preference of the opponent can be made based on her previous bids (we denote this class by $A(G)$). These attributes will be assigned a coefficient R_j in the "Guess Coefficients component (as described in Section 2.2.5).
- Attributes for which no reliable information about the preference weights of the opponent can be made from his previous bids (denoted by $A(NG)$). These attributes are assigned a default value $R_j = 2$, which is empirically chosen between the values for attributes for which there is an indication they are important to the opponent (from her past bids) and those attributes which are less important to her (as will be explained in Sect. 2.2.5).

After establishing the value of the R_j parameters, the estimation of the actual weights for the attributes $j \in A_{unknown}$ is computed as:

$$W_j = \frac{R_j}{\sum_{k \in A_{unknown}} R_k} \sum_{k \in A_{known}} W_k \quad (2.9)$$

It is possible that no reliable information can be obtained from the opponent's past bids for any of the attributes. Then all distribution coefficients will be equal and applying the above formula results in equal distribution of the remaining weight between private attributes, or formally: $A_{unknown} = A(NG) \implies \forall j, k \in A_{unknown}, W_j = W_k$.

2.2.5 Guess Coefficients Component

This component analyses the opponent's bids and, for those attributes for which a trend is reliable detected, returns a value for the remaining weight distribution coefficient R_j .

In the current model the explicit assumption used in guessing (for the Seller's side only) is that, everything else being equal, a human Buyer would prefer a better quality item to a poorer quality one. Otherwise stated there exists a (partial) ordering of the attribute values such as: $evaluation(good) > evaluation(fairlygood) > evaluation(standard) > evaluation(meager) > evaluation(none)$. We define the *Attribute Value Distance* $AVD(j)$ for each attribute $j \in A$ as the distance between values for an attribute in two successive bids, on an ordinal scale. For example, given the above ordering, the distance between good and fairly good is 1, while the difference between good and standard is 2. It is important to note that this attribute value distance does not depend on the exact values the opponent assigns to these labels - since in the current model this information is private (not disclosed to the other). From the experiments run for this domain, we observed that this simple ordering information can lead to a reasonably good heuristic. Partial ordering information is usually sufficient to make a good prediction about the opponent's preferences in the negotiation (i.e. if this distance is known only for some labels, this is enough).

Next we need a mapping of the detected concession distances to the remaining weight distribution coefficients introduced in Section 2.2.4 (see Table 2.2). The values for the above coefficients were determined experimentally as follows: first between each two different labels (representing quality levels) an initial value was computed by subtracting their distance value from 4 (the maximum distance). Then the parameters were adjusted to provide a best fit for the results over a large number of tests. It is important to note that the mapping in Table 2.2 should be seen as domain-specific. It led to good experimental results for the model and tests reported in this chapter (as we will show next). However, if this bilateral negotiation model is adapted for other domains (for example with more attributes, different number of levels/attribute etc.), then probably another choice of parameters would be needed.

Another issue to be discussed is how many successive bids in the negotiation trace need to be analyzed in order to make a prediction for R_j . From our empirical tests we observed that in most cases it is sufficient to adjust the R_j parameter based only on the first 3-5 bids. This can be explained by the fact that, in our model, being partially cooperative, agreement over the attributes with discrete values occurs in the first rounds of the negotiation and usually the last rounds can be characterized as "haggling" over the only continuous attribute, the price.

$AttributeValueDistance(j)$	R_j
0	6
1	4
2	3
3	1
4	0.5

Figure 2.2: Remaining Weight Distribution Coefficients assigned to Attribute Value Distances for all attributes $\in A(G)$

2.3 Implementation & Experimental validation

The model introduced in Section 2.2 was implemented in the DESIRE agent platform [31]. However, the conceptual negotiation model presented in Section 2.2 is platform independent. In fact, after the publication of our original research [115, 118], a negotiation model which re-uses our mathematical model, but implemented in the more commercial Java Aglets platform was presented in [201].

In these tests, the model is assumed as symmetrical, i.e. both buyer and seller use the concession and opponent modeling heuristics described in Sect. 2.2.1-2.2.4. An important difference is that only the seller agent can use the guessing heuristic. This choice is explained by the fact that the model must be seen as partially cooperative and, furthermore, in most e-commerce scenarios, it is the seller that has more information to make such an estimation.

This section first discusses the experimental set-up used in testing the model presented in Section 2.2. Following, a full example trace is presented for the implementation. Finally the aggregate experimental results (for different test parameters) are presented.

2.3.1 Experimental set-up

In order to test the robustness of the above model, we considered the following dimensions:

- The number of attribute weights revealed (i.e. the degree of “openness” of the negotiation)
- Whether guessing is used or not by the seller
- The choice for the attribute importance factors
- The evaluations for the attribute value levels

The importance factors assigned to different attributes are presented in Table 2.3. Note that these were also chosen to be typical for the space of possible preferences of users in this domain. The values presented in Table 2.3 are raw importance factors, which are then normalized to add up to 1, using the formulas from Section 2.2.1.

Next, we should check that these results hold for different possible value configurations. Again the search space here is very large, so we must restrict our attention to a few profiles combinations, which are shown in Table 2.4.

	Tow hedge	Airco	Extra speakers	CD player	Price
Fully asymmetric	90 / 15	90 / 15	15 / 90	15 / 90	300 / 300
Partially asymmetric	53 / 53	90 / 15	15 / 90	53 / 53	300 / 300
Fully symmetric	53 / 53	53 / 53	53 / 53	53 / 53	300 / 300

Figure 2.3: Importance factors used for Buyer/Seller, for different levels of preference asymmetry

In our tests, we assume a business model in which the Seller prefers to sell the car for a standard price and not have to install extra accessories, but he is willing to do so in order to sell it. However, if he does have to install some accessories, as shown in Fig. 2.4, he would prefer them to be “standard” quality. On the seller side, these local utility values capture an ease utility for the seller to perform the required installation. For example, he may already have a stock of “standard” quality components ready to install, if asked, but very good components are difficult to get and must be ordered in advance. Furthermore, in real life, there may be guarantee issues in selecting these values.

On the buyer side, because in our model the values for the attributes represent quality labels, the distances between utilities assigned to each label can be interpreted as how “quality conscious” or selective that buyer is. For example, by looking at Table 2.4, a buyer of Profile 2 is more selective than a buyer of Profile 1, because his utility for “fairly good” and “standard” qualities drops quicker, when compared to the optimal quality level “good”. Other choices are possible, but in order to properly test the model the choice for the values must be asymmetrical - meaning the two parties would like different values for each attribute. Otherwise the parties quickly agree on the configuration (since their interests are convergent) and the negotiation reduces to haggling about the price.

	Profile 1	Profile 2
BUYER	100 / 85 / 70 / 30 / 0	100 / 70 / 50 / 35 / 0
SELLER	30 / 65 / 80 / 65 / 100	30 / 50 / 70 / 85 / 100

Figure 2.4: Value levels Good / Fairly Good / Standard / Meager / None for each of the 4 attributes

2.3.2 An example negotiation trace

In this section, we illustrate the model presented in Section 2.2 through an example. Here we take the negotiation between a Buyer and Seller with totally asymmetric preferences (see Table 2.3), where the only information revealed between parties is the normalized weight of 1 attribute (Tow hedge). For accessories, for both Buyer and Seller, profile 1 is used (see Table 2.4). For this example, we use the perspective of the Seller, which in our case is the party using guessing. For reasons of space, we can illustrate only a small part of the full

mathematical model, it should be enough for the reader to understand the rationale behind some of our design choices.

The attribute Tow Hedge has the following normalized preference weights (see Table 2.3): $W_{BUYER, TowHedge} = 90 / (90+90+15+15+300) = 0.1764$

$$W_{SELLER, TowHedge} = 15 / (90+90+15+15+300) = 0.0294$$

From the perspective of the Seller the preference Difference Coefficient for Tow Hedge will be:

$$\delta_{TowHedge} = (W_{Buyer, TH} - W_{Seller, TH}) / (W_{Buyer, TH} + W_{Seller, TH}) = (0.1764 - 0.0294) / (0.1764 + 0.0294) = 0.714.$$

A positive value close to 1 (as shown in 2.3), indicates this the attribute is more important to the other party (the Buyer). As the general tolerance (for the Seller side) in this case is $\tau_{gen} = 0.3$, the attribute specific tolerance will be $\tau_{TowHedge} = \tau_{gen} * (1 + \delta_{TowHedge}) = 0.3 * (1 + 0.714) = 0.514$. Since $\tau_{TowHedge} > \tau_{gen}$, a larger concession than average will be made towards the Buyer's requested value in this attribute. This can be seen in Table 2.7 as a large concession, in the first round from "none" to "fairly good". Next we exemplify the guessing of the opponent's weights discussed in Sections 2.4 and 2.5. We do this only after the first two rounds from the opponent's bids, though the mechanism is the same for subsequent rounds.

The Value Distances and Remaining Weight Distribution Coefficient, computed for the attributes of unknown weights are as follows (see Fig. 2.6 for the Buyer's first 2 bids and 1 for the coefficient mapping). For the attribute "Airco", the value distance is between levels "good" and "standard" (see trace in Fig. 2.6), thus 2, so based on the mapping in Table 2.6, we get $R_{Airco} = 3$.

Similarly, for the attributes "CD player" and "Extra speakers", the value distances are between "good" and "meager", thus 3 => $R_{CD_player} = R_{Speakers} = 1$.

Since $\sum_{j \in A_{unknown}} W_j = 1 - (15+300)/510 = 0.235$, the estimated weights are:

$$W_{AIRCO} = 3 / (1+1+3) * 0.235 = 0.141,$$

$$W_{CD_PLAYER} = 1 / (1+1+3) * 0.235 = 0.047$$

In this case, the estimations produced by the guessing are not far from the true (non-revealed) values of the Buyer: 0.176 for Airco and 0.0294 for CD player.

Tables 2.6 provides the complete trace of this negotiation from the perspective of the Buyer, while Table 2.7 does the same from that of the Seller. The vertical columns show the bids made by the two parties in successive rounds.

Figure 2.5 provides a visualization of the negotiation progress in the joint utility space (as automatically produced by the implementation in our software environment). For clarity, only the first 3 bids of the Buyer (marked with a 1(B), 2(B), respectively 3(B)) and the first 2 of the Seller (marked 1(S) and 2(S)) are shown. The remaining offers all lie close together in the straight line between point 3(B) of the buyer and point 2(S) of the seller.




Figure 2.5: Utility space corresponding to the example trace from Tables 2.6 and 2.7

BUYER	1	2	3	4	5	Closing bid
price	18000	17450	17968	18047	18083	18083
Tow hedge	good	fairly good	fairly good	fairly good	fairly good	fairly good
airco	good	standard	standard	standard	standard	standard
speakers	good	meager	none	none	none	none
CD player	good	meager	none	none	none	none
Utilities						
Own bid	1	0.9203	0.9130	0.9094	0.9068	0.9068
Seller's bid	0.7407	0.8782	0.8830	0.8864	0.8889	0.8889

Figure 2.6: The negotiation trace: BUYER's perspective

SELLER	1	2	3	4	5	6:accept
price	16900	18468	18404	18359	18325	18083
Tow hedge	none	fairly good	fairly good	fairly good	fairly good	fairly good
airco	none	standard	standard	standard	standard	standard
speakers	none	none	none	none	none	none
CD player	none	none	none	none	none	none
Utilities						
Own bid	1	0.9378	0.9296	0.9238	0.9195	0.8884
Seller's bid	0.3167	0.5932	0.8737	0.8838	0.8884	0.8884

Figure 2.7: The negotiation trace: SELLER's perspective

This is an interesting effect, which we have seen in a number of negotiation traces: after establishing mutually agreeable values for the discrete-value attributes, the agents seem to “walk” the Pareto-efficient frontier towards each other's bid. This corresponds to the hag-gling about the price from rounds 3-5 in Tables 2.6 and 2.7.


The effect is interesting, since neither one of the agents knows exactly where the Pareto-frontier lies. They cannot compute this information because they only have partial knowledge of the opponents utility function (in our experiments, the position of the frontier was computed after the fact using the full information, but only as a benchmark to measure performance – i.e. without giving this information to the agents). Since bargaining agents' heuristics only model the opponent preferences and do not try to actually predict where the frontier lies, the fact that they end up so close to it shows the proposed heuristics are working (of course, for this particular negotiation trace). This effect can be explained by the fact that, if opponent modeling is performed efficiently, the agents “discover” the Pareto-frontier implicitly, using their approximate opponent models.

2.3.3 Comparing traces from the same test set

We define a test set as the set of all negotiation traces which share the same Pareto-efficient frontier and therefore whose outcomes are directly comparable. Between the negotiations in the same set, the preferences of the two parties are the same: the only difference is the amount of information shared and their willingness to use guessing. Test sets are distinguished from each other by two (sets of) parameters:

- The level of asymmetry in the attribute importance factors (see Table 2.3)
- The distances between the value labels (cf. Table 2.4)

Comparing the efficiency of outcomes, both within the same test set and between test sets, is important in our setting, since, as is shown here, the efficiency of the outcomes which can be reached depends on how asymmetric the preferences of the parties are. First, we first discuss the results from a test set with maximum preference asymmetry (this setting also corresponds to the example in Sect. 2.3.2 above). Next, we discuss the results from two related test sets, but where preferences are more asymmetric. Finally, in Section 2.3.4 we present the aggregate results for all test sets, across all value distance settings and preference symmetry profiles considered.



Experimental setting	Price	Tow hedge	Airco	Extra speakers	CD player	Buyer Utility	Seller Utility
Closed negotiation(with or without guessing)	19018	Standard	Standard	Standard	Standard	0.865185	0.84177
One attribute weight revealed, without guessing	19178	Fairly good	Standard	Standard	Standard	0.86822	0.84054
One attribute weight revealed, with guessing	18325	Fairly good	Standard	None	None	0.906815	0.91946
Two attribute weights revealed, without guessing	18790	Fairly good	Standard	None	Standard	0.882741	0.87211
Two attribute weights revealed, with guessing	18325	Fairly good	Standard	None	None	0.906815	0.91946
Three attribute weights revealed (with or without guessing)	18428	Fairly good	Fairly good	None	None	0.914	0.90166

Distribution of final outcomes for the negotiations between a Buyer with a stronger preference for Drawing Hook and Airco, and a Seller with a stronger preference for CD player and Extra Speakers.

Figure 2.8: Outcomes for negotiations between a Buyer and Seller with asymmetric preference weights

Test set with maximal preference asymmetry

The table in Fig. 2.8 shows the final outcomes of negotiations involving a Buyer and Seller with asymmetric preferences and value profiles 1, while the graph in the same figure shows these outcomes are plotted w.r.t the Pareto-optimal frontier. The notation is: 1..3 denotes the number of attributes shared and NG/G denotes whether guessing is used or not. The Pareto frontier in the top graph from Fig. 2.8 is actually the same as the one in Fig. 2.5, just scaled between different values. In fact, the outcome reached in Fig. 2.5 appears as point 1G in Figure 2.8 (top). The irregular, non-convex shape of the Pareto-efficient frontier (computed according to [180]) is typical for real-life domains, where some attributes take discrete values and only some are continuous.

From the above test set we can already make some observations. First, more attribute weights shared improves the outcome, so the heuristic makes efficient use of incomplete preference information. In Table 2.8, this is illustrated by the fact that, as more attribute weights are shared, a better match is obtained between the preferences of the parties, with each one obtaining its preferred values in the more important attributes. Second observation is that the guessing heuristic may improve the outcome, sometimes considerably. In the trace presented for 1 or 2 attribute weights shared guessing helps bring the outcome very close to the Pareto-efficient frontier. For 0 attribute weights shared (i.e. perfectly closed negotiation), in this particular test set guessing does not help much (however there are test sets where it does). In the 3 attribute weights shared case the outcome without guessing is already Pareto-efficient. Note however that this case is not equivalent to fully open negotiation, because the evaluations for the values assigned to each quality level are still not revealed between parties.

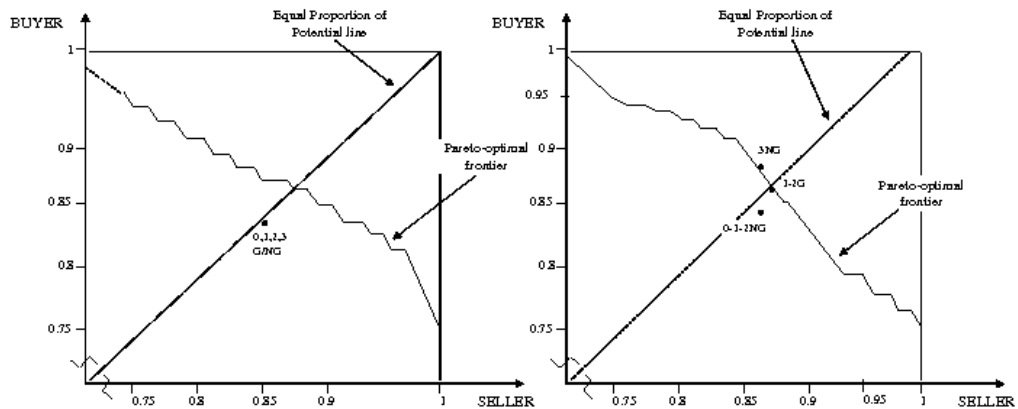


Figure 2.9: A (left side): Outcomes for a negotiation between a buyer and seller with completely symmetric preferences (i.e. all attributes have equal weights for both parties). B (right side) Outcomes of a negotiation between a buyer and seller with partially symmetric preferences (i.e. two attributes have equal weights, two are asymmetrical - c.f. Table 2.3

Test sets with symmetric and partially symmetric preferences

In this section we illustrate two other test cases, related to the one above. The distances between values for the buyer and seller are still 1 (cf. Table 2.4). Figure 2.9 plots the outcomes from two test sets: the first one in which the preference weights of both parties are the same across all 4 discrete-valued attributes, the second one in which only two attributes have equal preference weight, the other two having asymmetrical weights (see Table 2.3 for the exact values).

From Figure 2.9, it can be observed that, in fact, for more symmetric preferences revealing more information and/or using guessing does not make too much difference (the tables with the exact outcomes reached are not given here for lack of space, but they point to the same conclusion). In fact, for the case with completely symmetric preferences (Fig. 2.9(a)) we can see that all outcomes actually overlap. In this special case (equal weights across all attributes), the negotiation actually becomes a zero-sum game, since there are no mutually beneficial trade-offs between attributes, so the best that can be achieved is to settle on the middle of the range value.

2.3.4 Comparing results from all test sets

The above traces give a fairly good idea of the effects observed for different levels of preference asymmetry and openness. Subsequently, we tested all considered combinations of attribute level profiles (c.f. Table 2.4), number of weights revealed (c.f. Table 2.3), levels of openness regarding preference weights for different attributes (up to 4) and whether the agent chooses to make use of the guessing heuristic. Figure 2.10 shows the average utilities across tested profiles, grouped by the level of asymmetry in preference weight between parties. Within each group, from left to right the level of openness is varied from no attributes revealed and no guessing used to 3 attributes weights revealed.

Based on Figure 2.10, we can see that our observations from Section 2.3.3 generalize across profiles: both sharing more information and guessing improves the utility (on average). It can be seen that the more asymmetrical the preferences of the two parties are, the greater the scope for potential gains that can be obtained either by sharing more information or using the guessing heuristic. For example, for all profile combinations tested in the perfectly symmetrical preferences case, the outcome always had a 0% improvement, either from sharing more preference weight information or by using the guessing heuristic. By contrast in the partially symmetric preferences improvements were of the order of 3-4%, which went up to around 10% for asymmetric preference weights. This effect can be explained by the fact that our mechanism exploits precisely this preference asymmetry in order to increase the efficiency of the joint outcome for both parties. Another observation is that, if the overall concession speed parameter is roughly the same for both parties, the outcomes will always lie relatively close to the equal proportion of potential line, regardless of the guessing/openness model used. Otherwise stated, the overall concession for the bid level are similar, even though for each attribute may differ widely. This ensures that, if the negotia-




Figure 2.10: Average utilities for all profiles tested, for different cases of preference asymmetry and openness

tion outcome lies on, or close to, the Pareto-efficient frontier, it will also be relatively close to the Kalai-Smorodinsky bargaining outcome. This may be important, since some sources (e.g. [180]) consider closeness to this point as a measure of “fairness” of the negotiation outcome.

2.3.5 Human-computer experiments

The results reported in this Chapter refer only the automated negotiation case, i.e. the case when both Buyer and Seller are represented in the negotiation by automated software agents. The user (or owner of the agent) in this case, only needs to input its preference parameters that describe his/her utility function, and the software agent computes the bids/counter-offers on his/her behalf. We find it useful to mention, however, that other work performed using this negotiation system considered the case when humans propose their own bids against software agents. That work involved testing the negotiation model on 70 students, negotiating both against each other and against our automated agents. The full details and results of that work are outside the scope of this thesis (as the human tests were not performed directly by the author), but the interested reader can find a comprehensive description of the results in [22, 23].

Finally, a note should be made regarding the potential for exploitation of the system by other users. While the negotiation system proposed here is essentially cooperative in nature, the negotiation mechanism was designed to prevent obvious ways of cheating, like over-stating attribute preference weights. This is because each agent scales the sum of the preference weights declared by the opponent to 1. So an agent has no incentive to over-state his preferences for any attribute, since this may lead to the opponent making smaller or no concessions in other attributes. Furthermore, a system was added by which an agent stops negotiating when it detects insufficient concessions from the other in several successive bids, which should prevent situations where one party makes all the concessions. This made the system difficult to exploit in the human tests reported in [22, 23], but we note that these tests were made as “one run” (i.e. subjects did not get to learn to exploit the system through repeated interactions).

2.4 Discussion

In this section we provide an overview of existing work on negotiation that is most related to the work reported in this Chapter, such as to contrast them to this model. However, we leave a more complete review of negotiation techniques for Chapter 3, after the presentation of the other negotiation model discussed in the thesis.

In Gutman and Maes [91] a number of criteria and benefits are discussed of some different approaches to negotiation. This paper makes a strong argument for building partially co-operative negotiation systems, and shows that merchants often care less about profit on any given transaction and care more about long-term profitability, which implies customer satisfaction and long-term customer relationships. The work of Gutman and Maes served as a conceptual starting point for the negotiation model developed at the Free University of Amsterdam, which was further extended through the heuristics presented in this chapter. The ideas proposed in [91] are supported in our model by allowing consumers and provider agents to specify extensive multi-attribute profiles and degree of openness regarding preference information, and by developing heuristics that make use of this information to reach negotiation outcomes that aim to satisfy both parties.

Technically, the work that is most related to the one reported here is Faratin, Sierra & Jennings '03 [71]. Like Faratin et al., we start from the perspective of distributed negotiation, which eliminates the need of a central planner. As in [71], we also take the heuristic approach and we model agents that are able to jointly explore the space of possible outcomes with a limited (incomplete) information assumption. In [71], this is done through a trade-off mechanism, in which the agent selects the value of its next offer based on a similarity degree with previous bids of the opponent. In our design, we do not explicitly model trade-offs, yet the same effect is achieved through the asymmetric concessions mechanism. An advantage of our model over [71] is that we allow agents to take into account not only their own weights, but also those of the opponent in order to compute the next bid. In this way agents may exchange partial preference information for those attributes for which their owners feel this

does not violate their privacy. Also the initial domain information for the attributes with discrete (“qualitative”) evaluation is different. In [71], this consists of fuzzy value labels, while in our model it is a partial ordering of attribute weights.

Finally, a paper that provides a completely independent re-implementation of the model presented here is Shakshuki & Abu-Driaz’05 [201]. They propose a peer-to-peer, agent-mediated e-commerce system, that uses multi-issue negotiation as one of its core components. At the conceptual level, the negotiation model employed by [201] is basically the same as the one presented in this Chapter and was inspired by it, after the initial, conference publication of our work [117, 118]. Therefore, indirectly, this paper provides another validation of the conceptual negotiation model presented here, since it represents an independent re-implementation of this model in a different domain (web services) and using a completely different software platform (in [201], the Java Aglets platform was used instead of DESIRE).

2.5 Conclusions

The Chapter introduces a component-based generic agent architecture for partially cooperative multi-attribute negotiation. An application of the model is described in a prototype system for negotiation about cars, developed started from a practical application, partially in co-operation with, among others, Dutch Telecom KPN and Almende B.V. Rotterdam.

The main original contributions of this work are:

- It develops a generic component-based generic agent architecture that supports integrative multi-attribute negotiation, for both agent-agent and agent-human scenarios.
- Developing a heuristic that can use a limited amount of preference information to achieve outcomes on or close to the Pareto-efficient frontier.
- A guessing strategy that further improves the outcome of the negotiation, in case such partial information is not forthcoming from the negotiation partner.

The negotiation model described was implemented using the DESIRE software environment, and it follows component-based design principles of that framework. However, the applicability of the conceptual negotiation model presented in this chapter is not strictly limited to that framework. In fact, a negotiation framework that uses the heuristics developed here, but independently implemented after the publication of our initial work was presented in [201].

A possible limitation of this work may be the limited dimensionality of the domain considered, and in particular the fact that it implicitly assumes that bargaining agents’ utility functions over the multi-attribute space are linearly additive. This limitation has been dealt with in subsequent work by the author, presented in the next chapter.

Chapter 3

Modeling Complex Multi-Issue Negotiations Using Utility Graphs

3.1 Introduction

As shown in the introduction and the previous chapter, automated negotiation forms an important type of interaction in agent based systems for electronic commerce [147]. It allows buyers and sellers to determine the terms and content of a trade iteratively and bilaterally. Consequently, deals consisting of sets of complex goods or services can be tailored to the preferences of individual buyers and flexible to changing circumstances.

In this chapter, we consider the problem of a seller agent negotiating bilaterally with a customer about selecting a subset from a collection of goods or services, i.e. a bundle, together with a price for that bundle. Thus, the bundle configuration — an array of bits, representing the presence or absence of each of the shop's goods and services in the bundle — together with a price for the bundle, form the negotiation issues. Like the model presented in Chapter 2 (and the work of [71, 115, 126, 163, 207], among others), the techniques developed in this chapter aim to benefit from the so-called win-win opportunities. In the multiple item setting considered here, this means finding mutually beneficial alternative bundles of items during negotiations.

Several papers on automated approach the problem of finding win-win opportunities during bargaining through modeling the preferences of the negotiation partner (among them we mention Faratin et al. [71], Coehoorn & Jennings [163] and Jonker & Robu '04 [115], the last one serving as the basis of Chapter 2). However, an important limitation of all these approaches is they assume the issues under negotiation have *independent* (i.e. linearly additive) valuations for the negotiation partner. The approach taken by this chapter is also based on opponent modeling: we consider *interdependencies* between issues however, which, as we will show, makes the problem considerably harder. In order to model such

complex utility interdependencies between items, we introduce the novel concept of utility graphs.

Utility graphs build on the idea introduced in Chajewska and Koller [44] and Bacchus and Grove [9] that highly nonlinear utility functions, which are not decomposable in sub-utilities of individual items (such as in the seminal work of Raiffa [179]), may be decomposable in sub-utilities of clusters of inter-related items. They mirror the graphical models developed in (Bayesian) inference theory (cf. [140, 172]). Graphical models have been shown to be a powerful formalism for modeling decisions and preferences of other agents (see e.g. [30] for an overview). The idea behind using utility graphs in a multi-issue bargaining setting is to provide the seller with a formalism that can be used to efficiently explore the exponentially large bundle space. In this chapter, we show how utility graphs can be used to model an opponent's (i.e. customer's) preferences.

To illustrate our approach, consider the following application setting, which is quite natural from the point of view of current e-commerce practice. The seller represents an electronic merchant which aims to sell (subsets of) a set of items (e.g. books, CDs, pay-per view music tunes, news items etc.). He encounters and negotiates with random buyers who log in his website or web service at different times (thus all negotiations are bilateral and sequential). Buyer utilities are completely private (not known to the negotiation opponent), and the same holds for seller costs of providing the items. Finally, buyer preferences can be efficiently encoded in k -additive form (and, thus, are representable as utility graphs - c.f. definition in Sect. 3.3.2). Intuitively defined, k -additive utilities are a widely encountered and fully expressive class of non-linear utility functions, which be represented as a summation over partial functions over a set of [not necessarily disjoint] utility clusters, each containing up to k interdependent items (a full, formal definition is provided in Sect. 3.3.1). Although in this chapter we return to the e-commerce scenario to exemplify our approach, we note our model is not limited to an application setting: it applies to any setting where buyer preferences can be concisely represented in k -additive form.

We assume that, for privacy reasons, the seller is not allowed, between negotiations, to store preference information which is traceable back to *individual* buyers. This is actually a natural model for very large, open environments such as the Internet, where the same buyers and sellers are unlikely to encounter each other very frequently and most negotiation encounters are first time. Furthermore, there may be other reasons why personalized information cannot be stored in some settings (e.g. existing legislation). The seller can, however, store **aggregate** buyer preferences for the items he sells, information not traceable back to individual buyers.

Based on practice and existing literature, our negotiation approach (and resulting model) addresses a few important properties (which can be seen as desiderata of the model):

- The negotiation is maximally privacy preserving. Unlike combinatorial auctions, there is no **direct** preference information which either the buyer/seller needs to reveal to the opponent before the negotiation starts. An important reason for this (besides preserving privacy) is the difficulty the buyer may encounter in formulating his/her exact pref-

erences over the full set of items under negotiation. This is often a recurring situation in practice: the buyers do not have, or find difficult to formulate their exact valuation over all possible bundles. The number of possible contracts (i.e. bundles which can be formed) is exponential: for example, for only 50 items, there are $2^{50} > 10^{15}$ possible bundles. Even though they can be more compactly represented using a utility graph, it is, in general, it is far easier for a buyer to report her preferred combination at a given price level, rather than elicit the full preferences encoded in her utility graph (both in terms of structure and values). This argument resembles the perspective taken in preference elicitation with demand queries ([30, 32, 133]). This work identifies the need to elicit exact preferences as the main bottleneck encountered for the practical application of direct revelation mechanisms, such as combinatorial auctions. In our model, the only preference information which the agents (indirectly) reveal is that which can be gathered from the offers/counter-offers they exchange.

- The negotiation should maximize, as far as possible, the Pareto-optimality of agreed deals (formally defined, in Sect. 3.2, in terms of gains of trade associated to a contract).
- The number of negotiation rounds (i.e. offer/counter-offers exchanged during the negotiation) should be minimized as far as possible. This would enable our approach to be used in applications where time constraints or the impatience of buyers are limiting factors.

At the start of a negotiation process, the seller's approximation of the customer's utility graph represents some prior information about the *maximal structure* of the utility space to be explored. After every (counter) offer of the customer, this approximation is updated based on this offer. We show that, by using only a fairly weak assumption on the maximal structure of customers' utility functions, the updating procedure enables the seller to suggest offers that closely approximate Pareto efficiency. Moreover, efficient outcomes are reached after relatively few negotiation rounds, by efficiently exploiting the decomposable structure of complex utility functions.

The concept of the structure of a maximal graph of dependencies is a crucial part of our model. Intuitively defined, the maximal graph can be seen as a limitation of the types of nonlinear item-item dependencies which need to be considered when modeling the preferences of a randomly encountered buyer (i.e. mathematically equivalent to a limitation on the number of the terms of the polynomial describing a k -additive utility function, as formally presented in Sect. 3.3.2). In practical settings, this is a natural way to model the problem, since for most users their utility function over many binary issues or items, although it can contain non-linearities, is not arbitrarily complex. This maximal graph enables us to restrict the opponent modeling, from a super-exponential number of possible utility functions, to a more manageable number. This enables the seller to focus the search on the most promising region of the high-dimensional utility space. For this reason, the problem of finding a mutually efficient contract can be addressed in a limited number on negotiation steps, despite the high dimensionality of the contract space.

The issue of how the seller acquires this initial graph information is an important one

- and is addressed in the second part of our model. One method would be to elicit it from domain experts (i.e. an e-commerce merchant is likely to know which items are usually sold together or complimentary in value for the average buyer and which items are not). For example, if the electronic merchant is selling pay-per-item music tunes, the tunes from the same composer or performer can be potentially related - and this information can be used in negotiations with newly encountered customers.

We show this initial graph information can also be retrieved automatically, by using information from completed negotiation data. The implicit assumption we use here is that buyer preferences are in some way clustered, i.e. by looking at buyers that have shown interest (through their negotiation bids) for some combinations of items in the past, we can make a prediction about future buying patterns of a random customer. This holds, even if we assume we have not encountered a particular buyer in previous negotiations. In our model, only aggregate, not personalized preference data is maintained. Note that this assumption is not uncommon: it is a building block of most recommendation mechanisms deployed in Internet today [197]. In order to generate this initial structure of our utility graph, in this chapter we propose a technique inspired by collaborative filtering.

Collaborative filtering has been widely used in electronic commerce to provide buyer decision support, even when very large numbers of items are involved (e.g. Amazon.com offers over 2 million book titles). Recommendations are based on aggregate information obtained from a community or class of buyers, but they are not tailored specifically to the preferences of an individual buyer. By combining filtering with multi-issue negotiation a much greater degree of flexibility can be achieved, because, in addition to an estimation from aggregate preferences, deals can be tailored to the specific needs of the individual buyers. Therefore our work can be seen as establishing a link, in the form of a well-defined formal model (i.e. utility graphs), between customization mechanisms typically used in multi-issue negotiation and customization mechanisms used in collaborative/social preference filtering.

The remainder of this chapter is organized as follows. Section 3.2 presents the negotiation setting: it defines the efficiency criteria, the negotiation protocol and the top-level outline of the negotiation algorithm. Section 3.3 defines the concept of utility graphs, it describes how such graphs can be used to model complex utility functions and also introduces the concept of a maximal graph for a class of buyers. Section 3.4 defines the core of the negotiation model, by describing how utility graphs can be used for efficient exploration of the contract space and to learn opponent preferences in negotiation settings. Section 3.5 presents the method used in constructing the structure of the maximal utility graphs, based on concluded negotiation data. Section 3.6 presents the set-up and results from the experiments performed to validate the model. This Section consists of two main parts: validation of the negotiation model and validation of the graph retrieval algorithms. Section 3.7 presents a discussion of related work, in comparison to other approaches to this problem. Section 3.8 concludes the chapter, by highlighting the main contributions of our work and outlining some possible directions for further research. The chapter includes one appendix, containing a formal proof of equivalence between Pareto-efficiency and gains of trade maximization. This is incidental to the main objective of the chapter, but it is included for consistency.

3.2 The negotiation setting

We consider a buyer and seller who negotiate bilaterally over a set of n binary-valued issues or items, and one continuous issue, the price. Henceforth, we will refer to the binary-valued issues as items and to subsets formed with these items as bundles. Negotiations are conducted in an alternating exchange of offers and counter offers, using an alternating offers protocol. The offers and counter offers contain a n -dimensional vector of 0's and 1's representing an instantiation of the n items, plus a price offered/asked for the current bundle.

The utility (measured in terms of monetary value) a buyer assigns to any bundle of items is given by a non-linear (i.e. k -additive - as defined in Sect. 3.3.2) function that takes into account interdependencies between various items. The seller's utility for a bundle (measured in net monetary value) is the difference between the price received for a bundle and the costs incurred for providing a bundle. In the model used, the costs of the seller for providing the items are linearly additive: i.e. the bundle cost equals the cost of offering the items individually. Thus, in our model, the non-linearity in utility is mainly centered on the part of the buyer. However, both the buyer's utilities and the seller's cost represent private information, which remains undisclosed before or during the negotiation. Therefore, the negotiation setting can be described as double-sided incomplete information. For reasons outlined in the introduction, this is a natural requirement in bilateral, agent-mediated negotiation settings.

3.2.1 Net Utility functions of Buyer and Seller

Let $B = \{I_1, \dots, I_n\}$ denote the collection of n items a seller and buyer negotiate over. Each item I_i takes on either the value 0 or 1: 1 (0) means that the item is (not) purchased. Thus B has the domain $Dom(B) = \{0, 1\}^n$ (so there are 2^n possible bundles). The n -dimensional vector $\vec{b} \in Dom(B)$ denotes an instantiation of these n items. In our approach, the utilities assigned to different outcomes (combinations) are quasi-linear, i.e. represented by monetary units, rather than values normalized between 0 and 1, as in other negotiation models [71, 115, 163, 179, 232]. Quasi-linear utility functions have several advantageous properties (discussed in the next Section) and represent a natural choice in a wide variety of application settings that could be considered for bilateral negotiations, such as electronic commerce, task allocation, distributed logistics etc.

The utility function $u : Dom(B) \mapsto \mathbb{R}$ specifies the monetary value a buyer assigns to all (2^n) possible outcomes. Due to interdependencies between various items the function u can be highly nonlinear. The buyer's *net utility* (i.e. net monetary value) for purchasing a bundle \vec{b} for a price p , denoted by $nu_b(\vec{b}, p)$, is defined as follows:

$$nu_b(\vec{b}, p) = u(\vec{b}) - p \quad (3.1)$$

That is, $nu_b(\vec{b}, p)$ is the difference between the monetary value for acquiring (consuming) bundle \vec{b} minus the price p paid for purchasing bundle \vec{b} . The net monetary value of the seller

is computed as:

$$nu_s(\vec{b}, p) = p - Costs(\vec{b}) \quad (3.2)$$

Thus, the seller's net monetary value for the sales of a bundle \vec{b} for a price p is just the price minus the cost for selling the items. Currently, the seller has an *additive cost structure*: i.e., the bundle costs $Costs(\vec{b})$ equals the sum of the cost incurred when selling the items individually.

3.2.2 Using gains from trade as efficiency criteria

The most widely used performance criteria in multi-attribute negotiations is Pareto efficiency. Raiffa [179] provides a method to compute Pareto-efficient outcomes in the case utility functions of both parties are normalized between 0 and 1 (a choice made in several other multi-agent negotiation models e.g. [102, 115] etc.).

In our model, utilities are represented in monetary units instead of mappings between 0 and 1, or otherwise stated utility function are quasilinear. In the case of quasilinear utility functions, in order to determine Pareto-efficiency, it is enough to compute the gains from trade that can result from exchanging a certain bundle of items \vec{b} . The gains from trade are defined as:

$$GT(\vec{b}) = u(\vec{b}) - Cost(\vec{b}), \quad (3.3)$$

where $u(\vec{b})$ denotes the buyers monetary value for \vec{b} . The notion of gains from trade is well founded in the economic literature on trade (c.f. [88]). Moreover, for the above setting (where utility is expressed in monetary units) the set of bundles maximizing the gains from trade can be proven to be the same as the set of Pareto-efficient bundles (a formal proof is provided in Appendix 3.A of this chapter and [207]). Intuitively, the gains from trade can be seen as the maximal size of joint gains, which can be achieved through negotiation, while the continuous attribute, the price represents different ways to divide these joint gains.

3.2.3 Outline of the negotiation setting and protocol

The negotiation, in our model, follows an alternating offers protocol. At each negotiation step, each party (buyer/seller) makes an offer which contains an instantiation of 0/1 for all items in the negotiation set (denoting whether they are/are not included in the proposed bundle), as well as a price for this bundle. The decision process for each party, at each negotiation step, is composed of 3 inter-related parts: (1) take into account the previous offer made by the other party, (2) compute the contents (i.e. item configuration) of the next bundle to be proposed, and (3) compute the price to be proposed for this bundle.

In our model, the burden of exploring the huge bundle space and recommending jointly profitable solutions is passed to the seller, who must solve it by modeling the preferences of

the buyer. This is a reasonable model, in cases where an electronic merchant negotiates with different buyers in succession and tries to optimize both his own profits and buyer satisfaction (essential for building up a durable client relationship, which would generate repeated business). Furthermore, in many settings, one party can be seen as more knowledgeable than the other (e.g. in e-commerce domains it is reasonable to assume that electronic merchants are more knowledgeable than individual buyers, while in distributed logistics settings larger transportation providers have more knowledge of the market than random customers).

Following Gerding et al. [80] (among others), we view a multi-issue negotiation strategy as composed of two parts:

- The Pareto-search strategy, which enables agents to reach mutually profitable agreements, with incomplete or uncertain information about opponent preferences.
- The concession strategy employed during the negotiation, which can be one of the standard, time-dependent strategies, from the ones discussed in [70] (and other sources).

The focus of this chapter is on the Pareto-search aspect of the negotiation, in particular on proposing a search method for the case the utility functions of different agents are unknown and non-linear. The seller's approach is to search for a bundle that maximizes the gains from trade, since this increases the joint gains: the only way he can continue to offer a better concession on price, but also increase his own profits. Thus, there is a semi-cooperative nature of the negotiation: although the seller is selfish and tries to maximize his own utility, he achieves this by trying to find a better trade-off between issues, a trade-off which better matches the buyer's preferences, estimated from the buyer's responses so far. Such a mechanism also used in other negotiation approaches (e.g. [71, 163, 232]), although these only consider linearly additive utilities.

With respect to pricing, the seller starts the negotiation by posting an ask price for each item. This price reflects the maximum profit the seller expects to make by selling that item (i.e. the ask price reflects his maximal aspiration level). The lower bound of the seller's aspiration level (i.e. how far he is willing to go w.r.t. price concessions) is given by his vector of costs for providing each item, information which remains hidden from the buyer during the negotiation. The buyer also has a maximal aspiration level, which reflects the minimum price he expects to pay. The price the buyer will actually pay for a bundle depends on the content of that bundle and on the negotiation process itself. Both the buyer and the seller compute, for the current proposed bundle their current aspiration level and then make a time-dependent price concession with respect to that aspiration level. For consistency, results reported in this chapter refer to the *monotonic, time-dependent* concession case. However other time-dependent concession strategies usually employed in bilateral negotiations (such as hard-headed, Boulware [70]) have been considered.

It is important to note that the specific time-dependent concession strategy used refers only to prices, not to the search for a bundle configuration. With respect to the bundle configuration, both parties have an incentive to search for a Pareto-efficient combination, that maximizes the joint utility. For the seller side, this search is explicit, through opponent

modeling. For the buyer side, the bundles proposed simply reflect his/her preferences. In the current set-up, there is no obvious incentive for the buyer to falsely reveal his preferences, given the Pareto-search strategy of the seller, since revealing false preferences will lead to a bundle with sub-optimal Gains from Trade to be selected, which is also in its detriment. Because we make the explicit choice that the concession strategy only concerns the pricing of the bundles, it is also in the interest of the buyer to search for a good bundle combination, regardless of the price concession tactic he/she employs. Finally, we model time pressure and/or buyer impatience through a break-off probability: at each step there is a small risk of breakdown (a value of 2% was used in the simulations).

3.2.4 Assumptions about buyer knowledge

Given that our aim is to develop a model for handling incomplete preference information, an issue which needs to be discussed is the type of knowledge about the preferred combination the buyer can provide. The aim, in our model, is to handle buyers that do not (or are unable to) specify their preferences directly (i.e. cannot, or are not willing to, reveal the subutilities corresponding to the individual clusters/interdependencies, that form their utility function).

In our model, the buyer can maintain a vector of price expectations of items, based on the previous offers of the seller. With regards to the knowledge required, we assume only that the buyer is able to respond with (one of) the bundle combination(s) that, given his price expectations for the items at some point during the negotiation, maximizes his own utility. This is considerably less difficult for a buyer than revealing his exact preference function over all possible bundles. The assumption is similar to the assumption made in demand queries in preference elicitation literature (e.g. [133]), i.e. that the agents are required to only locally reveal their preferred combination, for a given price level. The negotiation problem, is however, more complex than combinatorial preference elicitation with a trusted proxy, since there are double sided-incomplete information and strategic considerations at play. Furthermore, as will be shown, our method is a heuristic rather than an exact elicitation method, which makes it computationally considerably more efficient.

3.2.5 Top-level negotiation algorithm used by the seller

Alg. 1 gives an outline of the algorithm the seller uses in each negotiation step. There are several stages. First, the seller checks if the configuration of the last 2 offers of buyer and seller coincide and if the difference in price is below a given, small threshold δp . This threshold can be set as small as desired; in our current model is expressed as a percentage of the seller's last price. In case the difference is below the threshold, the seller accepts the buyer's last offer and the negotiation ends.

As briefly outlined in the introduction (and will be formally defined in the next Section), the seller models the preferences of his buyer in the form of a utility graph, in which each edge (cluster) is associated to a table of values. There are several main computational steps

during the negotiation: the opponent modeling step (line 2 of Alg. 1), in which the seller updates the values in the utility graph, based on the last offer made by the buyer, the counter-offer sent to the buyer is computed (line 4 of Alg. 1) and a step in which the price to be proposed next is determined (line 5 of Alg. 1). Note that there is a separation between the strategy the seller uses in opponent preference modeling and computing the best bundle and the price concession tactic.

Algorithm 1 Top level algorithm used by the seller

Denote by (\vec{b}_b, p_b) the previous offer of the buyer and by (\vec{b}_s, p_s) the previous offer of the seller.

1. If $\vec{b}_b = \vec{b}_s$ (configuration is agreed) and $p_s - p_b < \Delta p$ (difference in ask and offer prices is under some maximum acceptable threshold), then Success.
 2. Otherwise:
 3. Update the estimated utility graph of the buyer based on his past bid \vec{b}_{buyer}
 4. Compute (one of) the bundles \vec{b}^* with the highest gains from trade
 5. Compute the price to be proposed for \vec{b}^* such that it represents a linear time concession from my previous offer
 6. Propose this bundle and price to the buyer
-

3.3 Decomposable utility functions and their graphical representation

Recall that we consider a buyer who negotiates with a seller over a bundle of n items, denoted by $B = \{I_1, \dots, I_n\}$. Each item I_i takes on either the value 0 or 1: 1 (0) means that the item is (not) purchased. The utility function $u : \text{Dom}(B) \mapsto \mathbb{R}$ specifies the monetary value a buyer assigns to each of the 2^n possible bundles ($\text{Dom}(B) = \{0, 1\}^n$).

In traditional multi-attribute utility theory, u would be decomposable as the sum of utilities over the individual issues (items) [179]. In this chapter we relax this assumption by considering u decomposable in *sub-clusters* of individual items such that u is equal to the sum of the *sub-utilities* of different clusters.

Definition 1 Let C be a set of clusters of items C_1, \dots, C_r (with $C_i \subseteq B$). We say that a utility function is factored according to C if there exists functions $u_i : \text{Dom}(C_i) \mapsto \mathbb{R}$ ($i = 1, \dots, r$ and $\text{Dom}(C_i) = \{0, 1\}^{|C_i|}$) such that $u(\vec{b}) = \sum_i u_i(\vec{c}_i)$ where \vec{b} is the assignment to the variables in B and \vec{c}_i is the assignment to the variables in C_i , induced from the assignment \vec{b} . We call the functions u_i sub-utility functions.

The factorization of a decomposable utility function is not unique. In this chapter, we use the following factorization, which is a relatively natural choice within the context of negotiation. Single-item clusters ($|C_i| = 1$) represent the individual value of purchasing an

item, regardless of whether other items are present in the same bundle. Clusters with more than one element ($|C_i| > 1$) represent the *synergy effect* of buying two or more items; these synergy effects are positive for complementary items and negative for substitutable ones.

Note that the synergy value associated with a cluster (set) of k items represents an additional value, beyond the utility associated with any of its proper subsets. This method of defining utility functions is, in fact, equivalent to what existing literature refers to as a k -additive form.

3.3.1 The k -additive utility form

The k -additive form represents an important class of representing decomposable utility functions described in existing literature (see [46, 55]) (the same class is referred to as the “polynomial representation” in some sources, e.g. [133]).

For unbounded k , the k -additive form is fully expressive, meaning any utility function over a set of (binary) issues (or bundles of items) can be represented in k -additive form (cf. [46]). In practice, however, the maximum size of k can nearly always be bounded to a small value. For instance, if we denote by x_1, \dots, x_n the instantiation of the set of n items, the expression for a 4-additive utility form (i.e. taking a maximum $k=4$) is given by:

$$U(x_1, \dots, x_n) = \sum_{1 \leq i \leq n} \alpha_i x_i + \sum_{1 \leq i, j \leq n} \alpha_{i,j} x_i x_j + \sum_{1 \leq i, j, k \leq n} \alpha_{i,j,k} x_i x_j x_k + \sum_{1 \leq i, j, k, p \leq n} \alpha_{i,j,k,p} x_i x_j x_k x_p \quad (3.4)$$

In the binary case, x_1, \dots, x_n represents a vector of 0 and 1, denoting whether an item is (or is not) considered in the combination being evaluated, the reals $\alpha_1 \dots \alpha_{n,n,n,n}$ are the parameters of the function, while the maximum k , henceforth denoted by k_{max} (same k as in “ k -additivity”) is the maximum rank of the polynomial, i.e. all the polynomial terms having a rank above k_{max} have the coefficients $\alpha = 0$. For example, linearly additive functions form a subclass of k -additive class, where $k_{max} = 1$. The expression given in Eq. 3.4 corresponds to a 4-additive utility function.

There is a natural connection between the definition in Def. 1 and Eq. 3.4. Each cluster corresponds a term of the polynomial, the number of clusters (also denoted by r) represents the number of terms in the polynomial expression with $\alpha \neq 0$. Only terms corresponding to clusters in the graphical model constructed for the buyer will be considered in the negotiation algorithm.

In practice, in the experimental part of this chapter, we limit our experiments for bilateral negotiation to $k_{max} = 4$. This means that, in the negotiation tests we performed, the polynomials modeling the utility functions of buyers are assumed to have maximum degree bounded to 4 (c.f. Eq. 3.4). Furthermore, the collaborative filtering method we propose,

which is used by the seller to get an approximation of the super-graph of buyer's utilities, was developed for a $k_{max} = 2$.

Intuitively, this means that for the case some structural information is available about the structure of the utility preference of the buyer (i.e. which terms of the polynomial are most likely to appear in his utility function in Eq. 3.4), we have validated the approach for up to 4th degree polynomials. The entire approach, including collaborative filtering, which assumes no personalized information exists about the utility function of any particular buyer agent, was investigated for 2nd degree polynomials.

Although the type of non-linearity our negotiation approach handles is restricted, from the point of view of the potential application domains, it is reasonable to assume a bounded k_{max} . For example, if our generic negotiation model would be applied business-to-consumer (B2C) e-commerce, then it is reasonable to assume buyers would have difficulty in assessing the additional benefits of very high-order synergies (this is also a reasonable assumption in many B2B e-commerce scenarios). Moreover, because the terms of the polynomials (corresponding to clusters in the graph) overlap, the problem remains significantly more complex than in the linear dependencies case.

3.3.2 Using graphs to model complex utility functions

The factorization defined above can be represented as an undirected graph $G = (V, E)$, where the vertexes V represent the set of items I under negotiation. An edge between two vertexes (items) $i, j \in V$ is present in this graph if and only if there is some cluster C_k that contains both I_i and I_j . More generally, if we allow for more than binary dependencies, we can define *hyper-edges* between k items. A hyper-edge between any k items: $i_1, \dots, i_k \in V$ is present in this graph iff there is some cluster C_k that contains all items I_{i_1}, \dots, I_{i_k} . We will henceforth call such a graph G a *utility graph*¹.

Example 1 Let $B = \{I_1, I_2, I_3, I_4\}$ and $C = \{\{I_1\}, \{I_2\}, \{I_1, I_2\}, \{I_2, I_3\}, \{I_2, I_4\}\}$ such that u_i is the sub-utility function associated with cluster i ($i = 1, \dots, 5$). Then the utility of purchasing, for instance, items I_1, I_2 , and I_3 (i.e., $\vec{b} = (1, 1, 1, 0)$) can be computed as follows: $u((1, 1, 1, 0)) = u_1(1) + u_2(1) + u_3((1, 1)) + u_4((1, 1))$, where we use the fact that $u_5((1, 0)) = 0$ (synergy effect only occur when two or more items are purchased). The utility graph of this factorization is depicted in Fig. 3.1. There exist (indirect) interdependencies between all 4 goods, because the choice of I_1 is influenced by the relations between I_2, I_3 and I_2, I_4 .

To give a numerical example, suppose: $u_1(1) = \$7$, $u_2(1) = \$5$, $u_3((1, 1)) = -\$5$, $u_4((1, 1)) = \$4$, $u_5((1, 1)) = \$4$. Moreover, all item costs are equal to \$2: i.e., $c(I_1) = c(I_2) = c(I_3) = c(I_4) = \2 . In this case the bundle with the maximum gains from trade (i.e. the bundle denoted by \vec{b}^* in Alg. 1) is: $(0, 1, 1, 1)$, which has the net monetary value

¹The concept of cluster defined for utility graphs corresponds to the concept of cliques in probabilistic networks. However, for consistency, we use only the term "cluster" throughout this chapter.

of $\$5 + \$4 + \$4 - 3 * \$2 = \$7$. From this simple example we can already highlight an important problem. The assignments for items I_1 and I_3, I_4 influence each other indirectly, through the assignment for I_2 (although there are no direct links from I_1 to I_3, I_4). This feature is exploited by our decomposition algorithm.




Figure 3.1: The utility graph that corresponds to the factorization according to C in Example 1. The $+$ and $-$ signs on the edges indicate whether the synergy effect are positive or negative.

At the computational level, each cluster is represented by a *joint utility table*, which assigns a utility value for all combinations of instantiations with 0/1 of items in that cluster (this is similar in concept to the joint probability tables, used to represent inter-dependent variables in probabilistic networks). Although (as mentioned above), in the experiments, we limit the maximum size of the cluster to a small k_{max} (which means up to k_{max} items are *directly* interdependent in utility, i.e. linked by an edge, any number of items can be indirectly interdependent, depending on the connectivity of the graph. Any two issues whose corresponding vertexes are connected by a path in the utility graph are, potentially, interdependent. For example, the factorization corresponding to the graph discussed in Example 1 and Fig. 3.1 from above has $k = 2$ (only two-item dependencies), though all items are in fact interdependent through the choice of I_2 .

3.4 Negotiation heuristics based on utility graphs

As shown in Sect. 3, in our approach we model the buyer's utilities a graph. The structure of this graph, as well as the sub-utilities corresponding to different clusters constitute private information which is not revealed directly during the negotiation.

However, in our model, the seller does have some prior information to guide his opponent modeling. He starts the negotiation by having a super-graph of possible buyer graphs, i.e. a graph containing all possible inter-dependencies between the issues (items) which can be present in a given domain. The utility graphs of buyers form subgraphs of this graph. We call this graph a *maximal item inter-dependence graph*, because from the perspective of the seller, it gives the maximal set of clusters that need to be considered during the negotiation. Otherwise stated, the maximal graph tells the seller which terms of the k -additive sum in

Eq. 3.4 can potentially, have non-zero values for any random buyer he encounters in a negotiation.

The negotiation is still with double sided incomplete information, because: 1. The seller does not know anything about the values the buyer assigns to different issues (i.e. values corresponding to the clusters). These need to be estimated during the negotiation itself, from the counter-offers he makes. 2. The exact structure of the utility graph of the buyer is not known, the super-graph represents only a maximal (and potentially inaccurate) approximation.

The presence of this graph helps to greatly reduce the complexity of the search space on the side of the seller. The structural information contained in such a graph can be obtained either from a history of past negotiations or elicited from human experts. Note that in most domains it is reasonable to assume that the seller does know something about the goods he is selling. For example, if he is selling online pay-per-view journal articles, then articles within the same category (or with the same author) can be potentially related (though not guaranteed to be related for every buyer). A method for automatically constructing its structure is provided in Sect. 3.5.

Therefore, our model does not assume the seller has to know the exact structure of the utility graph of the buyer. For example, suppose two issues are assumed substitutable by the seller, so the utility of the combination containing both items is lower than the sums of utilities for individual items. If the buyer signals (through his bids) that he is willing to accept bundles containing both items, the seller will adjust the weight of this relation (i.e. adjust the values for the relevant cluster sub-utility) towards the sum of utilities for individual items. Conceptually, this is equivalent to removing the edge from the graph - or the corresponding term from the summation (which means the items are no longer assumed substitutable). Thus, in this model it is possible to conceptually remove edge dependencies (by adjusting their perceived weight towards 0), but it is not possible to add new dependencies - other than those present in the maximal graph. Thus, the maximal graph can be seen as a restriction on the utility space considered during the negotiation.

3.4.1 Selecting the best counter offer

The problem faced by the seller at each negotiation step (see Alg. 1, step 2) is to choose a bundle \vec{b}^* which has the highest gains from trade of the 2^n bundles. More formally stated: $\vec{b}^* = \operatorname{argmax}_{\vec{b}}(\hat{GT}(\vec{b})) = \operatorname{argmax}_{\vec{b}}(\hat{u}(\vec{b}) - \operatorname{Cost}(\vec{b}))$. In the above equation and subsequent ones we use the notation \hat{GT} and \hat{u} , to distinguish them from the actual values GT and u , since the seller does not know the *true* utility values of the buyer. He only has an estimation of these values, which are updated after receiving a (counter) offer.

A straight-forward, “brute-force” solution to determine \vec{b}^* is to generate all bundles \vec{b} and select one which has the highest (estimated) gains from trade. Since this involves 2^n steps at *every* iteration, it is clearly not feasible for large n , so a heuristic is needed to reduce this search space. Now, suppose the utility graph is decomposable into two or more completely disjoint parts (no overlapping vertexes). We can then compute an optimal sub-bundle for

each of the parts and merge them. However, for the more general case we still need a method for reducing the complexity of the search, when the original graph (or a large sub-component of it) is not decomposable in such a straight-forward way. We do this by applying ideas from graph decomposition theory to the utility graph $\hat{G} = (\hat{V}, \hat{E})$.

Informally, the decomposition of a graph is a family of small, not necessarily disjoint, subgraphs $\hat{G}_1, \dots, \hat{G}_k$ (for some $k \in \mathbb{N}$), the union of which makes up the initial graph. Associated with a decomposition $\hat{G}_1, \dots, \hat{G}_k$ is a collection of cutsets: a cutset is a subset of vertexes that belong simultaneously to the same two subgraphs. (See [61, 215] for a more formal discussion of graph decomposition algorithms).

Example 2 Figure 3.2 depicts such a decomposition. The vertexes $\{Ia1, Ia2, Ia3, Ia4, Ic\}$ and $\{Ib1, Ib2, Ib3, Ic\}$ form the two cliques and corresponding subgraphs of \hat{G} . There are only two subgraphs therefore there is only 1 cutset; Ic forms this cutset because it is the only vertex that lies in both subgraphs.




Figure 3.2: Utility graph where the vertexes $\{Ia1, Ia2, Ia3, Ia4, Ic\}$ and $\{Ib1, Ib2, Ib3, Ic\}$ form the two cliques and corresponding subgraphs of \hat{G} . Moreover, $\{Ic\}$ forms the only cutset.

Given a decomposition $\hat{G}_1, \dots, \hat{G}_k$ of the graph G we can define $\hat{G}T_i : \text{Dom}(V_i) \mapsto \mathbb{R}$ as the predicted gains from trade that results from the sales of a subset of the item represented by the subgraph G_i . Here we use the formal notation $\text{Dom}(V_i)$ to represent the domain of possible values (in our case 0/1) assigned to the subset of vertexes V_i in subgraph G_i . Alg. 2 generates all possible combinations only for items that overlap between subgraphs (i.e. cutset items). Then, for all subgraphs it chooses the sub-combination that represents a *local* maximum for the gains from trade function in the considered subgraph, but subject to the constraint that the items that belong to more than one subgraph (i.e. cutset items) have the same values in all subgraphs. The best overall combination is chosen as a maximum of combinations of local maximums achieved for all possible instantiations for cutset vertexes (items).

Algorithm 2 Algorithm returns \vec{b}^* , a bundle with the highest gains from trade (i.e., $\vec{b}^* \in \operatorname{argmax}_{\vec{b} \in \operatorname{Dom}(B)} \hat{G}T(\vec{b})$)

Subgraphs $\hat{G}_1 = (V_1, E_1), \dots, \hat{G}_k = (V_k, E_k)$ and the union of all cutset nodes $S \subseteq V$ determined by the decomposition; $\hat{G}T_i : \operatorname{Dom}(V_i) \mapsto \mathbb{R}$ denotes the predicted gains from trade resulting from the sales of a subset of the items in subgraph G_i ; and $V_i[j], S[i] \in \{1, \dots, n\}$ denote the reference to the item in B that corresponds to the j^{th} and i^{th} vertexes in V_i and S , respectively.

1. $X := \emptyset$ // X will contain n -dimensional vectors
 2. For all $\vec{s} \in \operatorname{Dom}(S)$ {
 3. Initialize \vec{b} // \vec{b} is a n -dimensional vector
 4. For $(1 \leq i \leq k)$ {
 5. //Get local max. gains from trade in subgraph G_i consistent with \vec{s}
 6. $\vec{v}_i^* \in \operatorname{argmax}_{\vec{x} \in \operatorname{Dom}(V_i)} \hat{G}T_i(\vec{x})$
 7. s.t. $\vec{x}(l) = \vec{s}(m)$ if $V_i[l] = S[m]$
 for some $1 \leq l \leq |V_i|$ and $1 \leq m \leq |S|$
 8. For $(1 \leq j \leq |V_i|)$ $\vec{b}(V_i[j]) := \vec{v}_i^*(j)$
 9. } $X := X \cup \vec{b}$
 10. } return $\vec{b}^* \in \operatorname{argmax}_{\vec{x} \in X} \hat{G}T(\vec{x})$
-

Example 3 Consider the graph in Figure 3.2. Generating all bundle combinations and testing them takes 2^{p+c+q} steps. Our algorithm generates all possible combinations only for cutset C , then computes optimal sub-bundles for subgraphs A and B for each combination of C and merges them. This requires only $2^c(2^p + 2^q)$ steps. Since c in our case is of size 1 or 2, while p and q can be arbitrarily large, the decrease in the number of steps is exponential.

More generally, suppose the decomposition of a utility graph uses $k_S = |S|$ cutset vertexes and results in p subgraphs of sizes $|G_1|, |G_2|, \dots, |G_p|$ (where $|G_i|$ denotes the number of vertexes in subgraph G_i). The computational complexity of determining the optimal bundle (Alg. 2) can be written as:

$$O(2^{|S|}(2^{|G_1|} + 2^{|G_2|} + \dots + 2^{|G_p|})) \quad (3.5)$$

If we denote by $MaxV$ the number of vertexes in the resulting subgraph with the largest number of vertexes (i.e. $MaxV = \max_i |G_i|$), by $|S|$ the number of cutset vertexes (nodes) and by p the number of subgraphs G_i , then this complexity measure is upper bounded by the factor:

$$O(2^{|S|} 2^{MaxV} p) \quad (3.6)$$

An important problem to be addressed is how to find an (approximate) optimal partition (required by Alg. 2), given a utility graph structure, as well as the quality and computational complexity of this partition. This is the scope of Sect. 3.4.1.

Algorithms for efficient graph partitioning

Recall that Alg. 2 takes as input a graph decomposition that is already given and uses it to compute the optimal bundle combination. In this Section, we look at the problem of how can we partition any random graph structure, such as to assure a minimal computational cost for the search algorithm (Alg. 2). The upper bound for the complexity of Alg. 2 depends exponentially on two terms: $|S|$ (number of cutset or separator nodes) and $MaxV$, which is the size of the maximal subgraph. Thus, the partition required needs to be, as much as possible, a *balanced* partition: one that minimizes the number of vertexes in the largest subgraph.

The problem of efficiently determining a balanced graph partition has been extensively studied in OR and theoretical computer science literature, since it is a problem that appears in many application settings, thus is not specific to multi-issue negotiation. It can be formulated as a specific case of the *k-multiway separator problem* [69, 77]. To our knowledge, the most efficient approximate solution to this problem in existing literature is provided by Even et al. [69].

Formally, the problem can be stated as follows: given a graph $G(V, E)$, with n vertexes and m edges, where each vertex and edge is assigned a capacity (weight). In our case, we apply the partition to the unweighed case, thus we take the capacities: $w(v) = 1, \forall v \in V$ and $c(e) = 1, \forall e \in E$.

The problem is to partition the initial graph $G(V, E)$ into k parts, such that the sum of vertex weights in each of the resulting subgraphs is at most: $2 \frac{\sum_{v \in V} w(v)}{k}$. In order to preserve standard notation, we use k to denote the number of parts that the graph is partitioned into (denoted by p in Eq. 3.5 and 3.6). Thus k here has a completely different meaning than the k expressing the maximum degree of the polynomial, in k -additive utility. If n denotes the total number of vertices in the graph and $w(v) = 1, \forall v \in V$, this bound is equivalent to $\frac{2n}{k}$.

Thus, since each of the resulting subgraphs contains at most $\frac{2n}{k}$ nodes, thus we can write $MaxV = \frac{2n}{k}$. For a graph with all edge and vertex capacities equal to 1, [69] proposes an approximation algorithm where the total capacity of the separators (in our case denoted by k_S or $|S|$) is bounded by: $(2 + o(1))(\ln n)OPT_k$, where OPT_k is the optimal separator capacity that is possible for the given graph topology. The upper complexity bound for Alg. 2 from eq. 3.6 becomes:

$$2^{(2+o(1))(\ln n)OPT_k} 2^{\frac{2n}{k}} k \quad (3.7)$$

Where n is the number of vertexes (items) in the utility graph, k is the number of subgraphs the graph is to be partitioned into, $o(1)$ is a constant factor and OPT_k is the number of cutset nodes and represents a factor that is dependent on the actual topology of the graph.

Admittedly, the above equation is not a very tight bound - especially since it is still exponential in the factor OPT_k , which is a theoretical optimum, but depends on the graph topology and thus cannot be bound in advance. However, if we want to have a bound *for the general case graph* (i.e. without making any prior assumptions on our utility graph), the bound offered by k-balanced partition algorithm is probably the best we can do.

For many specific cases, the partitioning problem can be considerably more efficient, however. We cover this briefly, in the next Section.

Recursive decomposition of utility graphs. Restricted cases.

The complexity measure for Alg. 2, as discussed in Sect. 3.4.1 refers to one-step application of Algorithm 2. However, it is possible to apply Alg. 2 recursively on a utility graph G : first decompose G into m_1 subgraphs G_1, G_2, \dots, G_{m_1} , according to k_1 cutset nodes, then apply Alg. 2 again to find the optimal combination in each of the subgraphs G_1, G_2, \dots, G_{m_1} and then merge them for all combinations of cutset nodes k_1 . The advantage of this procedure is that the algorithm does not have to iterate over all cutset nodes $k_1, k_{1,2} \dots k_{1,m_1}, k_{2,1,m_2} \dots k_{2,m_1,m_2}, \dots$, but the iteration over the cutset nodes is localized in each subgraph. This can achieve considerable computational savings, but iff. the graph structure allows for a suitable recursive decomposition.

To illustrate, for trees, it is natural to apply Alg. 2 as follows: the root of the tree is also taken as the cutset node. For all instantiations of the root, the local optimal combinations corresponding to all the items (vertexes) in the subtrees are computed and then merged. The local optimal combinations in each subtree can then be computed recursively, by applying the decomposition and search algorithm in each subtree. Thus, it is straightforward to show that recursive decomposition works on trees and graphs of (limited) bounded treewidth (up to 2-3), which can be approximated to trees. For the general case, however, obtaining an approximate partition through the minimum b-vertex separator algorithm (c.f. Sect. 3.4.1) and then applying Alg. 2 of the resulting partition is the optimal way to proceed, given available state-of-the-art in graph decomposition algorithms.

Finally, we should mention that, in practice, for the experiments performed in this chapter, we randomly generated and tested several graph structures (including some complex ones, to assure robustness of the results). However, the method we employed to generate random structures was such as to always assure their decomposability, thus we implicitly set k and OPT_k through the graph generation method (see Sect. 3.6.3 for a discussion of this issue).

3.4.2 Updating Sub-utility Functions

The search method described in Section 3.4.1 works on the utility graph \hat{G} , and corresponding sub-utility functions $\hat{u}_i(C_i)$ ($i = 1, \dots, |\hat{C}|$, where \hat{C} denotes the set of clusters). They represent the best model of the opponent (i.e. buyer) that the seller has so far. After receiving

a counter offer from the buyer, he will update this model. More precisely, he will update the sub-utility functions. In this Section, we will discuss the rule for updating these sub-utility functions.

The updating algorithm determines, for all $\hat{C}_i \in \hat{C}$, which combination of items in \hat{C}_i was asked by the buyer at the last iteration (where there are $2^{|\hat{C}_i|}$ possible combinations). Then it increases the expected monetary value of the buyer for that combination and degrades the other combinations in the cluster. Intuitively, the idea is to strengthen a link between vertexes (represented by the corresponding sub-utility value) whenever a buyer indeed expresses an interest in purchasing the items corresponding to the vertexes; otherwise the link is weakened. Algorithm 3 gives the actual updating rules.

Algorithm 3 Algorithm for updating sub-utility functions

The seller's and buyer's last bids contain the binary assignments \vec{b}_s and \vec{b}_b to the variables in B ; moreover $\vec{c}_{i,s}$ and $\vec{c}_{i,b}$ denote the assignments to the items (corresponding to nodes) in \hat{C}_i , induced by \vec{b}_s and \vec{b}_b (i.e. $\vec{c}_{i,s}$ and $\vec{c}_{i,b}$ are sub-arrays of \vec{b}_s and \vec{b}_b).

1. For $(1 \leq i \leq |\hat{C}|)$ {
 2. if $\vec{c}_{i,s} \neq \vec{c}_{i,b}$ {
 3. $u_i(\vec{c}_{i,b}) := u_i(\vec{c}_{i,b}) * (1 + \alpha_c)$
 4. For all $\vec{c} \in \{0, 1\}^{|\hat{C}_i|} \setminus \{c_{i,b}\}$
 5. $u_i(\vec{c}) := u_i(\vec{c}) * (1 - \alpha(i))$ }
-

Example 4 Suppose we have the cluster $C_i = \{I_3, I_5, I_6\}$ (for a $i \in \{1, \dots, |C|\}$). The buyer's last offer contains the combination $I_3 = 0$, $I_5 = 1$, and $I_6 = 1$. Then the expected buyer utility for purchasing item 5 and 6 is increased: i.e., $u_i((0, 1, 1)) = u_i((0, 1, 1)) * (1 + \alpha(i))$. The expected utilities for all other combinations in $\{0, 1\}^{|\hat{C}_i|}$ (namely $u_i(1, 1, 0)$, $u_i(1, 0, 1)$, $u_i(1, 0, 0)$, etc.) are decreased.

Parameter $\alpha(i)$ in Algorithm 3 defines how much weight should be given to the request from the buyer's last bid, in each cluster. A higher $\alpha(i)$ mean that the seller is more likely to give in to buyer's preferences for a cluster. A straightforward choice would be to assign the same α to all clusters. The seller would then only take into account the buyer's preferences. However, in our model, the seller also takes into account the expected gains from trade of a cluster, which also depend on his costs (unknown to the buyer). The rationale for this is in the way the constructed utility graph is used during the negotiation: it does not have to be a perfect model of buyer preferences, but it is used to search for the bundle with the highest Gains from Trade (which also depend on the seller's cost). We therefore define a factor called the Gains from Trade Importance Ratio (GTR) as:

$$GTR(i) = \frac{GT_i(\vec{c}_{i,b})}{GT(\vec{b}_b)} \quad (3.8)$$

for all $i = 1, \dots, |C|$, the clusters obtained by the decomposition from the previous Section. In the above equation, \vec{b}_b denotes the buyer's last offer, and $c_{i,b}$ denote the binary sub-vectors of \vec{b}_b , containing the instantiations with 0/1 of the items contained in each cluster C_i .

Intuitively, this ratio provides a measure of the cluster's importance weight, by comparing the gains which can be obtained in this cluster to the gains for the whole bundle. The measure can be compared to the inverse of total number of clusters $1/|C|$, for the purpose of fine-tuning the alpha parameter for each cluster (although the clusters vary in size in the conducted experiments from 1-4 items, the above ratio still provides a useful heuristic). The cluster-specific α is then computed as the sum between a fixed and variable component (computed by a sigmoid function): i.e.,

$$\alpha(i) = \alpha_{fixed} + \alpha_{var} * \frac{1}{1 + e^{\beta * (GTR(i) - 1/|C|)}} \quad (3.9)$$

for all $i = 1, \dots, |C|$. Here the parameter β is a positive value, which gives a measure of how steep the sigmoid function is. After conducting a number of experiments, we observed that transforming the function into a simpler step function (by assigning $\beta \rightarrow \infty$) still works well in many experimental tests. The values assigned to the alphas were determined empirically for each set of experiments.²

The rationale behind the above formula is the following: if a cluster has a high importance for the seller (i.e. if $GTR(i) > 1/|C|$), then the concession made for this particular cluster will be small (equal to α_{fixed}). Intuitively, this means that for this cluster, the costs of the seller are low, so the seller should keep insisting more on offering his own values for the items in this cluster for longer, since he knows he can offer them cheaper (the buyer does not know this, because he does not know the cost structure of the seller). For clusters with relatively low gains from trade (i.e. if $GTR(i) < 1/|C|$), there is not much difference between the offer of the buyer and that of the seller - therefore the seller can agree to the values asked by the buyer in that cluster without much perceived utility loss.

Note that our learning rule entails that in all clusters at least a small, non-negative adjustment α_{fixed} is made towards the buyer's offer. We made this choice since it assures that our algorithm guarantees convergence (i.e configuration agreement) within a limited number of negotiation steps and, thus, the negotiation is guaranteed to converge to an agreement at some point (although how efficient this agreement is depends on the tuning of parameters and needs to be experimentally investigated).

It is important to note that the aim of our learning algorithm is not to exactly learn the buyer's preferences or converge to his/her true utility graph model. Rather, it is a heuristic, that aims to find a good outcome in a complex utility space under uncertainty, for the multi-issue negotiation setting we consider. This approach is somewhat similar to other algorithms developed in the preference elicitation literature (e.g. [133]), but the aim in multi-issue negotiation research is different, since the goal is not to converge to an exact opponent model.

²To exemplify, common parameter settings were: $\alpha_{fixed} = 0.1$ and $\alpha_{var} = 0.3$.

Exact preference elicitation is not strictly required in a bilateral negotiation setting: our goal is to obtain a model of the buyer that is sufficiently good to find a jointly agreeable bundle (given the cost vector of the seller), and not to learn exactly all the utility clusters' values that describe the buyer's utility function. Indeed, for the rather complex utility functions we consider in our experimental analysis, the job of learning the exact buyer values could not be solved in an average of 50 negotiation steps (as shown in Sect. 3.6). But for the multi-issue negotiation problem, this is not needed: we do not need an exact model of buyer preferences, only one good enough to compute and propose a jointly agreeable contract.

3.5 Constructing the structure of utility graphs using aggregate negotiation data

The previous sections have focused on the problem of modeling the negotiation process itself, with an individual buyer, starting from the structure of a maximal preference graph, which encode dependencies to be considered. We have speculated how this initial preference information could be obtained, in a given domain. For example, it may be the case that an electronic merchant knows which of the items he is selling are potentially related (e.g. music tunes from the same album, related web services, books covering the same topic etc.). In this section, we show that it can also be learned automatically, from previous aggregate negotiation data, through a technique inspired from work on collaborative filtering. The section is structured as follows. In Sections 3.5.1 and 3.5.2, we present an overview of collaborative filtering and how it relates to our negotiation model. In Section 3.5.3 we formally define the structural graph for a set of buyers, which is our target criteria, while Sections 3.5.4 and 3.5.5 describe how similarity measures are obtained from previous negotiations data. Finally, Sections 3.5.6 and 3.5.7 propose a method for building the structure of the resulting graph, using these similarities.

3.5.1 Collaborative filtering: brief introduction

Collaborative filtering [197] is the main underlying technique used to enable personalization and buyer decision aid in today's e-commerce, and has proven very successful both in research and practice. The main idea of collaborative filtering is to output recommendations to buyers, based on the buying patterns detected from buyers in previous buy instances. There are two approaches to this problem. The first of these is use of the preference database to discover, for each buyer, a neighborhood of other buyers who, historically, had similar preferences to the current one. This method has the disadvantage that it requires storing a lot of personalized information and is not scalable (see [197]). The second method, of more relevant to our approach, is **item-based collaborative filtering**. Item based techniques first analyze the user-item matrix (i.e. a matrix which relates the users to the items they have expressed interest in buying), in order to identify relationships between different items, and then use these to compute recommendations to the users [197].

3.5.2 Overview of our filtering and negotiation approach

There are two main stages of our approach (see also Figure 3.3):

1. Using information from previously concluded negotiations to construct the structure of the utility super-graph. In this phase the information used (past negotiation data) refers to a class of buyers and is not traceable to individuals. This is the focus of Sect. 3.5.
2. The actual negotiation, in which the seller, starting from a super-graph for a class (population) of buyers, will negotiate with an individual buyer, drawn at random from the buyer population above. In this case, learning occurs based on the buyer's previous bids during the negotiation, so information is buyer-specific. However, this learning at this stage is guided by the structure of the super-graph extracted in the first phase. The algorithms used in this stage were covered in Section 3.4.

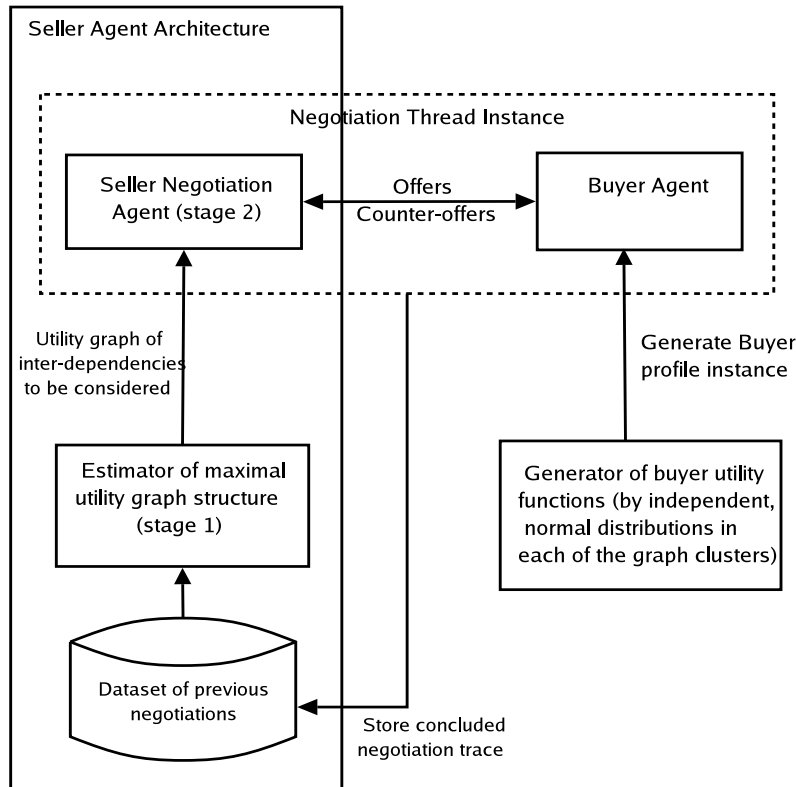


Figure 3.3: Top-level view of our agent architecture and simulation model

Phase 2 is described in our previous work [186]. The rest of this chapter will focus on describing the first phase of our model, namely retrieving the structure of the utility super-graph from previous data.

3.5.3 Minimal super-graph for a class of buyers

The definition of utility graphs given in Section 3.3 corresponds to the modeling the utility function of an individual buyer. In this chapter, we call the utility graph of an individual buyer the *underlying* or *true* graph (to distinguish it from the *retrieved* or *learned* graph, reconstructed through our method). The underlying graph of any buyer remains hidden from the seller throughout the negotiation.

We do assume, however, that the buyers which negotiate with a given electronic merchant belong to a certain class or population of buyers. This means the utility buyers assign to different bundles of items follow a certain structure, specific to a buying domain (an assumption also used indirectly in [79, 197, 207]). Buyers from the same population (or class) can be expected to have largely overlapping graphs, meaning there is a subset of dependencies that are much more likely to appear than others for a random buyer in that class. This does not entail that all buyers will have all dependencies specific to that class.

To illustrate this concept, consider the example of an e-commerce merchant negotiating over the configuration of a bundle of n goods with a random buyer. For simplicity, let's assume the buyer's utility graph can contain only binary dependencies (i.e. his utility function is 2-additive). Even so, there are n^2 possible dependencies (corresponding to edges) in the utility graph of the buyer (e.g. 2500 for 50 binary issues, 10000 for 100 binary issues etc.). The number of possible *utility graph structures* that a random buyer can potentially have is the size of the power set of the number of dependencies: $P(n^2) = 2^{n^2}$, a quantity which is so large that the problem of efficiently capturing buyer preferences seems to require solving an intractable modeling problem.

In practice, however, the seller knows that not all n^2 dependencies are relevant, and considering a much smaller set of dependencies usually leads to extracting most gains from trade from buyer preferences. Otherwise stated, the super-graph of buyer preferences to be considered is likely to be much sparser than fully connected: in fact most dependencies will not occur in the function of any buyer. In the following, we define these intuitions more formally.

Definition 2 Let $A = \{A_1, \dots, A_n\}$ be a set (class, population) of n buyers. Each buyer $i = 1..n$ has a utility function u_i , which can be factored according to a set of clusters $C_i = \{C_{i,1}, C_{i,2}, \dots, C_{i,r(i)}\}$. We define the super-set of clusters for the class of buyers $A = \{A_1, \dots, A_n\}$ as: $C_A = C_1 \cup C_2 \cup \dots \cup C_n$.

In graph-theoretic terms (as shown in Section 3.3), the set of clusters C_i according to which the utility of a buyer A_i is structured is represented by a utility graph G_i , where each

binary cluster from $\{C_{i,1}, \dots, C_{I,r(i)}\}$ represents a dependency or an edge in the graph. The super-set of buyer clusters C_A can also be represented by a graph G_A , which is the minimal super-graph of graphs G_i , $i = 1..n$. In practice, our negotiation algorithm uses a graph structure that is reasonably close to the minimal one, but may contain spurious (extra) or even missing edges.

$C_1 = \{\{I_1\}, \{I_2\}, \{I_3\}, \{I_1, I_2\}, \{I_2, I_3\}, \{I_2, I_4\}\}$. This super-graph is minimal, because if we were to add the dependency $\{I_1, I_3\}$ to C_A we would also obtain a super-graph, though not the minimal one.

It is important to note that the above definition for the utility super-graph for a class of buyer refers only to the **structure** (i.e. clusters C_i) and makes no assumption about the sub-utility **values** (i.e. functions u_i) in these clusters. To illustrate the difference, suppose that at a structural level, there is a complementarity effect between two items. However, for some buyers in the population, the utility value corresponding to this dependency may be very high (i.e. it is important for the agent to get both items), while for others it is much more moderate (or even close to zero).

3.5.4 Extracting information from concluded negotiation data

Suppose the seller starts by having a dataset with information about previous concluded negotiations. This dataset may contain complete negotiation traces for different buyers, or we may choose, in order to minimize bias due to uneven-length negotiations, to consider only one record per negotiation. In our model, we take the first bid a buyer makes in each negotiation thread. An alternative would be to take the final outcome of the negotiation - however, as this results after a process of negotiation and matching with the seller costs, it is not as reflective of buyer preferences as his/her initial asking offer.

The considered dataset is not personalized, i.e. the data which is collected online cannot be traced back to individual customers (this is a reasonable assumption in e-commerce, where storing a large amount of personalized information may harm customer privacy). However, in constructing of the minimal utility graph which the customers use, we implicitly assume that customers' preference functions are related - i.e. their corresponding utility graphs, have a (partially) overlapping structure.

Our goal is to retrieve the **minimal super-graph** of utility interdependencies which can be present for the class or population of buyers from which the negotiation data was generated.

We assume that past data can be represented as a $N * n$ matrix, where N is the number of previous negotiation instances considered (e.g. up to 3000 in the tests reported in this chapter) and n is the number of items (e.g. 50 for our tests). All the data is binary (i.e. with values of "1" in the case the buyer asked for this item or "0" if he does not). In the remainder of this Section, we use the following notations:

- N for the total number of previous negotiation outcomes considered

- For each item $i=1..n$, $N_i(1)$ and $N_i(0)$ represent the number of times the item i was (respectively was not) asked by the buyer, from the total of N previous negotiations
- For each pair of items $i, j = 1..n$ we denote by $N_{i,j}(0,0)$, $N_{i,j}(0,1)$, $N_{i,j}(1,0)$ and $N_{i,j}(1,1)$ all possibilities of joint acquisition (or non acquisition) of items i and j .

From the above definitions, the following property results immediately: $N_{i,j}(0,0) + N_{i,j}(0,1) + N_{i,j}(1,0) + N_{i,j}(1,1) = N_i(0) + N_i(1) = N_j(0) + N_j(1) = N$, for all items $i, j = 1..n$.

3.5.5 Computing the similarity matrices

Item-based collaborative filtering [197] works by computing "similarity measures" between pairs of items. The literature on item-based collaborative filtering uses two main criteria for computing similarities between pairs of items: **cosine-based** and **correlation-based** similarity. Both had to be adapted for our specific problem, i.e. retrieving the utility super-graph for a class of buyers from previous negotiation data. In particular, we needed to derive the expressions for the binary case, since existing mathematical definitions [197] are given only for real-valued preference scores.

In the following, we present the resulting expressions in both cases in separate subsections. As we will later show in the experimental part, only correlation-based similarity was found to work for our task, i.e. retrieving utility graphs from past data, but for completeness sake, in this chapter, we will report the formulas and experimental results we performed for both similarity criteria.

Cosine-based similarity

Cosine-based similarity is only useful for detecting complementarity effects between pairs of items. The resulting item-item similarity matrix contains only positive entries (between 0 and 1), a higher number denoting a stronger potential similarity. The formula to compute the entries in the is:

$$Sim_{compl}(i, j) = \frac{N_{i,j}(1,1)}{\sqrt{N_i(1) * N_j(1)}} \quad (3.10)$$

Correlation-based similarity

For correlation-based similarity, the resulting similarity matrix contains both positive and negative values (between -1 and 1). We first we define for each item $i = 1..n$, the average buy rate:

$$Av_i = \frac{N_i(1)}{N} \quad (3.11)$$

The following two terms are defined:

$$\psi_1 = \begin{aligned} & N_{i,j}(0,0) * Av_i * Av_j - N_{i,j}(0,1) * Av_i * (1 - Av_j) \\ & - N_{i,j}(1,0) * (1 - Av_i) * Av_j + N_{i,j}(1,1) * (1 - Av_i) * (1 - Av_j) \end{aligned}$$

and the normalization factor:

$$\psi_2 = \sqrt{\frac{N_i(0) * N_i(1)}{N}} * \sqrt{\frac{N_j(0) * N_j(1)}{N}}$$

The values in the correlation-based similarity matrix are then computed as:

$$Sim(i, j) = \frac{\psi_1}{\psi_2} \quad (3.12)$$

3.5.6 Building the super-graph of buyer utilities

After constructing the similarity matrices, the next step is to use this information to build the utility super-graph for the class of buyers likely to be encountered in future negotiations. This amounts to deciding which of the item-item relationships from the similarity matrices should be included in this graph. For both similarity measures, higher values (i.e. closer to 1) represent stronger potential complementarity. For substitutability detection, the cosine similarity uses a different matrix, while the correlation-based it is enough to select values closer to -1.

Ideally, all the inter-dependencies corresponding to the arcs in the graph representing the true underlying preferences of the buyer should feature among the highest (respectively the lowest) values in the retrieved correlation tables. When an interdependency is returned that was not actually in the true graph, we call this is an excess (extra, erroneous) arc or interdependency. Due to noise in the data, it is unavoidable that a number of such excess arcs are returned. For example, if item I_1 has a complimentary value with I_2 and I_2 is substitutable with I_3 , it may be that items I_1 and I_3 often do not appear together, so the algorithm detects a substitutability relationship between them, which is in fact erroneous.

The question on the part of the seller is: how many dependencies should be considered from the ones with highest correlation, as returned by the filtering algorithm? There are two aspects that affect this cut-off decision:

- If too few dependencies are considered, then it is very likely that some dependencies (edges) that are in the true underlying graph of the buyer will be missed. This means that the seller will ignore some interdependencies in the negotiation stage completely, which can adversely affect the Pareto-efficiency of the reached agreements.
- If too many dependencies are considered, then the initial starting super-graph of the seller will be considerably more dense than the “true” underlying graph of the buyer

(i.e. it contains many excess or extra edges). Actually, this is always the case to some degree, and in [186] we claim that Pareto-efficient agreements can be reached starting from a super-graph of the buyer graphs. However, this super-graph cannot be of unlimited size. For example, starting from a graph close to full connectivity (i.e. with n^2 edges for a graph with n issues or vertexes) would be equivalent to providing no prior information to guide the negotiation process.

In the general case, we consider graphs whose number of edges (or dependencies) is a linear in the number of items (issues) in the negotiation set. Otherwise stated, we restrict our attention to graphs in which the number of edges considered is some linear factor k times the number of items (vertexes) negotiated on. Framed in this way, the problem becomes of choosing the optimal value for parameter k (henceforth denoted by k_{opt}). The choice of the parameter k_{opt} (and, implicitly, of the density of the maximal utility graph to be considered) should depend on the negotiation stage of our model - since it is the negotiation stage that actually makes use of this structure. The following section presents the method used for choosing this parameter.

3.5.7 Minimization of expected loss in Gains from Trade as cut-off criteria

Denote by $N_{missing}$ the number of edges that are in the “true”, hidden utility graph of the buyer, but will not be present in the super-graph built through collaborative filtering. Similarly, we denote by N_{extra} the number of excess (or erroneous) edges, that will be retrieved, but are not in the true utility graph of the buyer. (in the experimental analysis in Section 3.6.6, rather than working with absolute figures, we find it more intuitive to report these measures as percentages with respect to the number of edges in the true underlying utility graph of the buyer).

The number of edges which are missing (not accurately retrieved) or excess (too many extra edges) depends on the accuracy of the underlying collaborative filtering process. More precisely stated, the number of missing edges depends on 3 parameters: the type of filtering used (correlation or cosine-based), the amount of concluded negotiation records available for filtering (we denote this number by N_r) and the number of edges considered in the cut-off criteria, k . Formally, we can thus write: $N_{missing}(corr, N_r, k)$. In this section we focus, however, exclusively on choosing a value for k , and consider the other two parameters as already chosen at the earlier step. Thus we simplify the notation to: $N_{missing}(k)$ and $N_{extra}(k)$, respectively.

As discussed in Section 3.5.6, both having missing and too many extra edges influences the efficiency of outcome of the subsequent negotiation process. Our goal is to choose a value for k that minimizes this expected efficiency loss during the negotiation. The efficiency loss, in our case, is measured as the difference in Gains from Trade which can be achieved using a larger/smaller graph, compared to the Gains from Trade which can be achieved by using the

“true” underlying utility graph of the buyer (in earlier work [186, 207], we have shown that maximizing the Gains from Trade in this setting is equivalent to reaching Pareto optimality).

In order to estimate this error rate, we consider a second negotiation test set, different from the one used for filtering. The purpose of this second test set is to obtain an estimation of the loss in gains from trade which occurs if we use a sparser/denser graph than the true underlying graph of the buyer. This is done by removing and/or adding random edges to the utility graph the seller starts the negotiation with and measuring the effect by re-running a number of negotiation threads. The ideal value for the size of the super-graph will be the one that minimizes the expected loss in Gains from Trade, compared to the gains from trade which can be achieved by an optimal outcome (henceforth we use the abbrev. “GT loss”).

In more formal terms, the expected utility loss for using k edges can be written as:

$$E_{loss_GT}(k) = \max(E_{loss_GT}(N_{missing}(k)), E_{loss_GT}(N_{extra}(k))) \quad (3.13)$$

Thus, for each value of the number of cutoff edges k , we compute the loss expected as maximum between the losses of having too few and too many edges. As we will show in the experimental results (Section 3.6.6), a too small value of k may lead to many edges being missed, compared to the true utility graph of the buyer. A too large value of k may lead to considering very dense graphs to start with, which also damages the negotiation process, albeit more gently (as will be shown in Section 3.6.6).

Thus we denote the optimal choice of k as:

$$k_{opt} = \operatorname{argmin}_k E_{loss_GT}(k) \quad (3.14)$$

Note that our criteria for choosing k in Eq. 3.14 relies on a similar intuition as “min-max regret” decision criteria, sometimes used in preference elicitation problems involving multiple issues [32, 39].

The application of the above principle to selecting a value of k_{opt} in practice will be discussed in Section 3.6.6, after we introduce the experimental results we obtained for different cut-off criteria.

3.6 Experimental Analysis

This section covers the results from the computer experiments performed to test the negotiation and retrieval models formally presented in Sections 4 and 5. Following the structure of the chapter itself, this section is divided in two main parts. In Subsections 6.1 to 6.3 we describe the results from experiments performed on the negotiation model itself (i.e. the algorithms presented in Section 4). In Subsections 6.4 to 6.6 we provide results from the experiments for the collaborative graph retrieval methods (presented in Section 5).

3.6.1 Experimental analysis of the automated negotiation heuristics

There are two main dimensions, which we want to test in our model:

- The distance to Pareto-efficiency of the outcome reached (measured, in our case, as the average percentage from the maximum possible gains from trade)
- The time taken to reach a solution, measured as the number of negotiation steps.

However, the preference structures, especially of the buyer side, are considerably complex, thus we need to ensure that the results are robust enough to hold in a wide range of possible settings. In particular, there are two sets of parameters which can be varied:

- The parameters of the distributions used to generate values for each of the clusters of buyer's utility graph and seller's costs, for a given graph topology are generated at random, according to a set of normal probability distributions.
- The structure (topology) of the Buyer's utility graph itself (thus not only the utility values per cluster) can also be generated at random, based on several classes of graph topologies. This is needed in order to show that the results are not just specifically fitted for one graph structure. Moreover, it enables us to study how the type of topology considered affects the performance.

To intuitively explain the difference, we can consider the k -additive form of the utility function which, in our model, is represented in graphical form (see Equation 3.4, Section 3.3.1). In this case the distributions of cluster utilities represent the numerical values which will be added up in these sums, while the topology of the graph represents which combinations of terms will appear in this summation (implicitly, all terms for which there is no edge in the graph, have the corresponding term equal to zero).

The organization of this Section is as follows. First we present the a restricted case, in which we vary the dispersions of the distributions of the sub-utilities assigned to different clusters, for one graphical structure comprised of 20, respectively 50 vertices. These are presented in Section 3.6.2. Section 3.6.3 presents the full simulation results, by examining how well these results generalize across different, randomly-generated graph structures, for graphs comprised of 50 vertices.

Generic experimental set-up for negotiation tests

Our negotiation model was tested in negotiations of different dimensions, involving 10, 20 and 50 binary-valued issues. Most of the tests and analysis reported in this chapter refers to the case when utility graphs of the buyer contains 50 vertices (i.e. binary issues) and 80 binary dependencies, 50 of which represent positive pairwise synergies and 30 negative ones. This enables us to explore a wide range of graph topologies in our tests, the choice of these

values being founded in the expected properties of the random graphs, as will be explained in Sect. 6.4. Graph structures (i.e. topologies) were generated at random, but following a certain distribution of edges into sub-graphs or topics (cf. Sect. 3.6.3).

However we also tested many other configurations of smaller/larger graphs, as well as configurations involving higher-order synergies between up to 4 items (i.e. hyper-edges connecting up to 4 vertexes), with consistently good results. For lack of space and need of consistency, we cannot cover all the tests for all dimensions here, but to provide the reader some idea of the results for the case of smaller graph dimensions, for comparison, we include some results for the 20 binary issues case as well.

Both buyer monetary utilities, in each cluster (i.e. interdependency) and seller costs are generated from normal distributions. On the seller side, the cost of each of the n items is normally distributed according to $N(\mu_{cost}, \sigma_{cost})$. On the buyer side, the value of each of n individual items is normally distributed according to $N(\mu_{gains}, \sigma_{gains})$. Moreover, the synergy effects between *subsets* in each of the above 80 binary clusters is normally distributed according to $N(0, 2\sigma_{syn})$.

To somewhat limit the number of parameters, we set $\sigma = \sigma_{costs} = \sigma_{gains} = \sigma_{syn}$. The parameter σ captures the problem of finding Pareto-efficient solution very, nicely: the higher σ the higher the likelihood of complementarity and substitutability effects, hence the higher the likelihood of non-linearity, in the problem.

The mean of the cost distribution for each item, on the seller's side, μ_{costs} , is set to 1. For the mean of the distributions for the buyer μ_{gains} , we conduct experiment with 3 different values: 1.1, 1.25, 1.5. Otherwise stated, the valuations of the Buyer are, *on average* 10%, 25% to 50% greater than those of the Seller. In the reported tests, σ takes 8 values, ranging from 0 to 5. In other words, we consider to whole spectrum from no randomness, and consequently linear preferences, to a high degree of randomness, and consequently (with probability) highly nonlinear preferences. (i.e. a normal distribution with a mean of 1 and spread 0.5 is rather centered around the mean, while a normal with mean 1 and spreads of 4 or 5 is virtually "flat", thus sums over such distributions are more non-linear and difficult to distinguish).

For each combinations of values of μ_{costs} , μ_{gains} , and σ , 100 tests were performed. Therefore for each result point in the graphs from Fig. 4 and 6 refer to averages over 100 tests. All of these settings for generating cluster values were applied to different randomly-generated graph structures, leading to a large number of negotiation tests.

In all of the tests reported in this chapter, we considered a buyer and the seller who use a *time-dependent, monotonic concession strategy* in proposing prices the bundles of goods. Other concession strategies were implemented with reasonably good results, but we acknowledge that there exist conceivable strategies (e.g. hard-headed), that do not lead to similar results or, in some cases, do not enable agents to reach a negotiation agreement at all.

As shown in Sect. 3.2.4, the scope of this chapter is to provide a robust model that enables agents to reach agreements in negotiations with complex utility functions and with incomplete information. Otherwise stated, our focus is on the Pareto-search aspect of the

negotiation, for which, in the complex setting considered, some implicit degree of cooperation is needed. Regarding price concessions, we selected a widely used strategy from the literature, which is suitable for our setting, leaving the exhaustive comparison of the effects of other concession strategies to future work.

In the following, we structure the presentation as follows. First, in Section 3.6.2 we report and compare the results for two graph structures, involving 20 and 50 binary issues. Next, in Section 3.6.3 we take the higher dimension case from our set-up (i.e. the 50 vertex case) and we examine the influence of different topologies (i.e. structures) of utility graphs on the negotiation process. First, we describe our method for generating different random graph structures to encode buyer preferences and we compare the experimental results obtained from each of these.

3.6.2 Experimental results for differently sized graphs

In this Section, we present the results of negotiation, using two different graph sizes. Results for the case involving 20 issues are presented in Fig. 3.4, while results for the case of 50 issues are presented in Fig. 3.5. The basic experimental settings are the same as presented in the experimental setup, i.e. Sect. 3.6.1 above. Each point was produced from 100 negotiations and the error bars give the resulting variance.

In both cases, the right-side figures highlight the results with respect to reaching an agreement over the bundle configuration (i.e., the actual content of the bundle). By agreement, we mean that thereafter the bundle content no longer changes. After such an agreement is reached, it may take more negotiation rounds before bargainers agree upon the final price. The bundle to be traded, and thus the Pareto-efficiency of a deal — which is the focus of this chapter — is then already determined, even if the parties continue to haggle about its final price.

These results are for just one structure (topology) of the utility graphs, as follows. For the 50 issue graph, we have 80 binary dependencies (edges in the graph, as described above) and we use the "Ring" topology (see Sect. 3.6.3 below). For the 20-issue case, we have used a graph with 24 dependencies, but 4 of these are 4-way dependencies and 8 are 3-way dependencies (representing hyper-edges in the graph).

From examining the results from Figs. 3.6.1 and 3.5 already some (partial) conclusions can be highlighted:

- Regarding the ability to find bundles close to maximum efficiency, the figures show that our approach performs well, for both graph dimensions. For the 20 issue graph, the approach is able to extract on average 97% of the maximal Gains from Trade, while for the 50-issue configuration only slightly less at 94% of the optimal GT. Worthy of note is the performance, in both cases for the higher values of σ (e.g., σ between 4-5). For these values, (with all likelihood) the problem becomes highly nonlinear and the buyer's utilities will show a high variance. Thus, for large values of sigma

the normal distributions become virtually flat, thus the mixture of such distributions becomes harder to detect.

- Regarding the number of steps required to reach agreement on the configuration, we do see a significant increase in this number for larger values of σ . Furthermore, here we also see a significant difference (more than a 2-fold increase) between the 20 and 50-issue cases. Clearly, for higher-dimensional problems and larger values of σ , the negotiation becomes more difficult to solve, so more steps are needed to correctly update the model of buyer's preferences and to find a good bundle. Nevertheless, a bundle very close to maximal efficiency is usually found, even in these more difficult cases.

Of course, the above conclusions were reached by varying the cluster potentials, but not the graph structures themselves. In order to fully validate our approach, we need to verify whether these results hold for a variety of different graph topologies - which correspond to different mixtures of utilities obtained from the distributions in each cluster.

3.6.3 Set-up and results for different graph structures

As mentioned in Section 3.6.1, we tested our models for graphs with $n = 50$ binary issues and $m = 80$ random clusters (edges). The question we try to answer here is: given these coordinates, how are graph structures generated? One obvious solution would be to generate the graph completely at random - i.e. select at random, for each of the 80 dependencies, the vertexes which they connect from the 50 vertexes/items (checking they do not coincide). However, for reasons explained in the sequel, here we use a more structured approach. In our approach, the n vertexes and m dependencies (edges) are first divided into k_t topics (subgraphs), not necessarily disjoint, each of them containing n/k_t vertexes and m/k_t dependencies (in the experiments reported in this chapter, we use $k_t = 5$, but other values have also been considered).³ Thus, each topic is a completely random subgraph, containing, in our case 10 (+/- 2) vertexes and 16 edges (differences are given by the fact that there are some vertexes that belong to more than one topic. As opposed to other uses of the term "topic" [47,48], in our usage of the word, topics are just a way to generate separately random subgraphs - topics are not separated from each other (except in one case: config. B from Fig. 3.6).

There are several reasons why we chose this form of generating random graphs:

- Having a graph which is decomposable, based on a set of cutset nodes, is important for the search algorithm. The size of the maximal subgraph (or topic) which results from such a decomposition has to be manageable, in order to be explored exhaustively. This property cannot be guaranteed for a random graph with a large number of edges

³As used in this chapter, the notion of "topics" is quite different from that of "clusters". Topics represent subgraphs, while utility clusters are represented by graph edges (or hyper-edges for more than 2-additive dependencies).

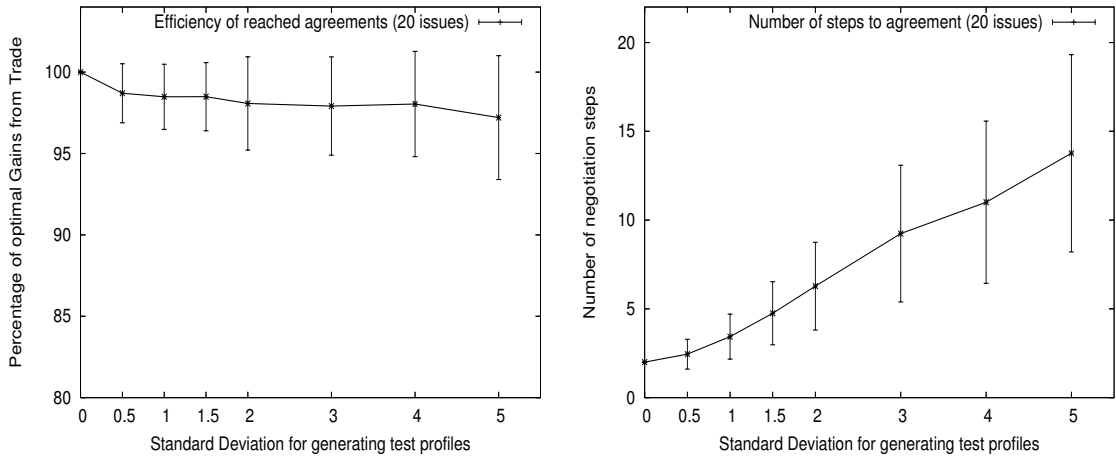


Figure 3.4: Percentage of the maximal Gains from Trade (compared to the Pareto-optimal bundle) and number of steps needed to configuration agreement, for a buyer utility graph with 20 issues.

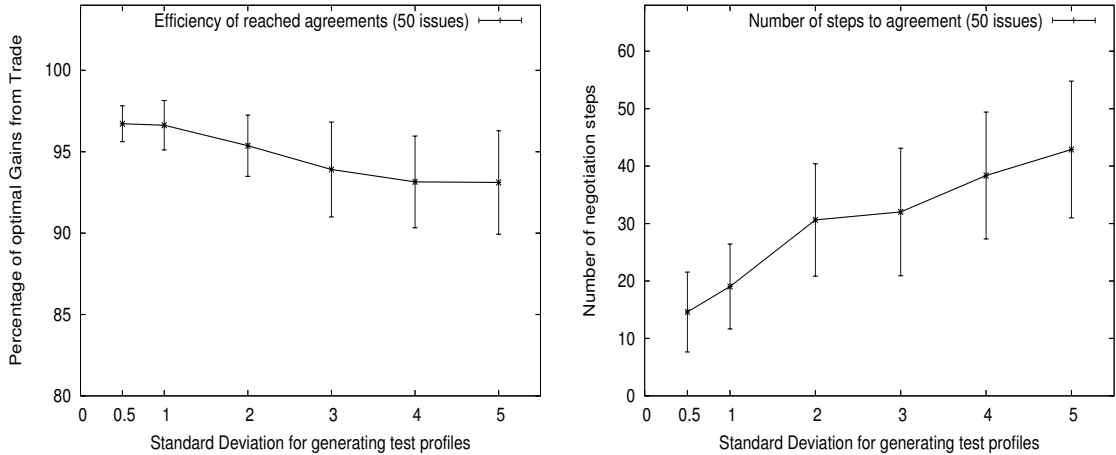


Figure 3.5: Percentage of the maximal Gains from Trade (compared to the Pareto-optimal bundle) and number of steps needed to configuration agreement, for a buyer utility graph with 50 issues.

[20, 191]. Of course the alternative is to choose much sparser graphs (with fewer edges). Nevertheless, we need relatively dense graphs (with a sufficient number of edges, representing linear inter-dependencies), in order to test the robustness of the model. This solution enables us to achieve both desired effects.

- The parameter k_t given above can provide a fine tuned-control for the type of graphs we generate. For example, for $k_t = 1$, we get back the fully random graph described above, while for $k_t = n$ (i.e. choosing one item per topic), we get back the linearly additive case (if there are no edges, hence no interdependencies). Thus, increasing k_t has the effect of increasing the structure of the resulting graphs. Nevertheless, graphs obtained with higher values of k_t are easier to decompose. We considered several values, but we argue $k_t = 5$, reported in this chapter provides a good trade-off between randomness and a decomposable structure.
- By using this procedure, we can consider different joinings of the random graph topics (which can share one or more cutset nodes). This is important, since the goal of our to obtain a characterization of how (and whether) different types of graph structures affect the negotiation algorithm results. Several cluster arrangements were selected, as shown in Fig. 3.6 - with the cutset nodes (i.e. nodes belonging to more than one topic) marked.

An issue to be discussed is the connectivity of the graph that results from this procedure. The issue of connectivity is important - since, in our case, the utility graph represents a number of interacting hyperplanes - which determines the complexity of the problem for the learning algorithms. In our case, we can show the following property:

Proposition 1 *The random graph structures generated from Fig. 3.6 (for configurations A, C, D and E) are connected with probability 97%.*

Proof: This is simply an application of random graph theory. For each random topic of 10 nodes and 16 randomly-generated edges, the probability that it is fully connected is equal to 99.4% (cf. [20], pg. 402 - for $c=2.25$). Since, for all configurations A, C, D, and E the topic are connected by construction (there are nodes belonging to more than one topic), the probability that the graph composed of 5 topics is connected is: $99.4^5 = 97\%$. This also explains our choice for 80 edges (clusters) for 50 issues.

In Fig. 3.7, the experimental results for all the graph topologies are shown. The experimental procedure is the same as described in Section 3.6.2, but repeated over all topologies. Each point from Fig. 3.7 (corresponding to a given topology/standard deviation for a cluster) represents an average taken over 100 tests. However, for the clarity and readability of the picture, we do not show the error bars here (they are around the same scale as those in Fig. 3.5 - around 5% for each data point).

Several conclusions can be drawn from the analysis of these figures. First, we observe that when the graph topology is loose (i.e. the graph is completely separated into 5 topics that are not connected to each other), the negotiation produces the best results and over 95% of the maximal Pareto-efficiency is achieved. The most “difficult” topologies for the algorithm were the “double chain” and the irregular ones, where there are several cutset nodes, belonging to more than 1 topic (in this case, the average efficiency achieved goes down to around 85%). However, it’s important to note that the variances of these results are




Figure 3.6: Quasi-random graph structures used in the experimental analysis. In order to describe them more intuitively, we use the names: A. Ring, B. Loose, C. Star, D. Double chain, E. Irregular shape1, F. Irregular shape2

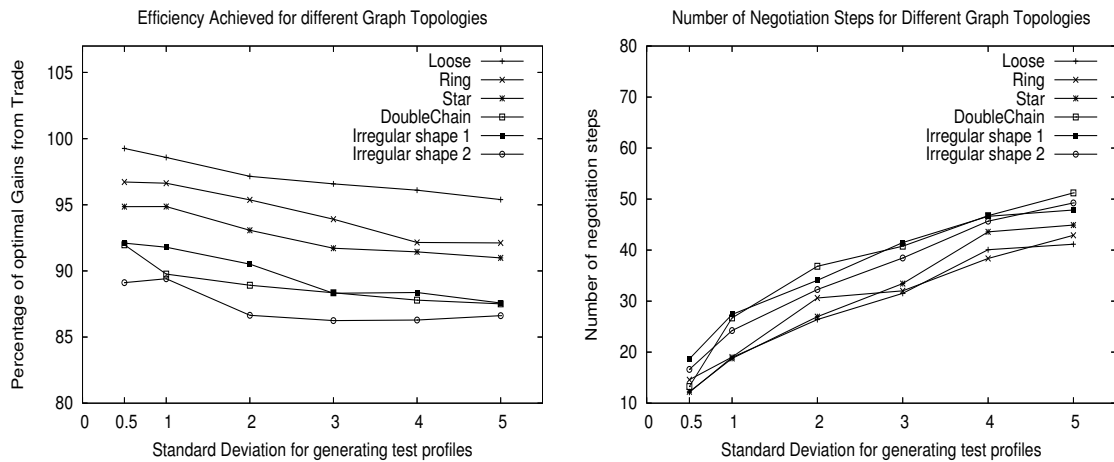


Figure 3.7: Percentage of the maximal Gains from Trade and number of steps needed to configuration agreement, for all graph topologies studied

also quite significant (around $\pm 5\%$ for each data point). Because of these high variances, the differences between the different graph structures (topologies) are, statistically, not very significant. These results show that the proposed negotiation heuristics can extract a high percentage of the gains from trade and are relatively robust to different graph structures (extracting 85% of the maximal Pareto-efficiency for complex topologies and high variance in cluster utility values -generated around a mean of 1 - is a robust result).

Regarding the number of negotiation steps required to reach the outcome, we observe a

similar effect: when the utility functions have a “loose” or “star” topology, negotiations take, on average, less steps than then in the “double chain” case, when there are more than 1 node in each cutset. However, we observe that overall, the standard deviation used in generating utilities in each cluster has a stronger impact than the topology of the graph itself. In terms of Fig. 3.7 b, the number of negotiation steps increases as we move towards higher standard deviations for all plots, by a higher factor than the differences between topologies.

However, the negotiation algorithm is able to extract over 85% of the Gains from Trade for all standard deviations and topologies, even if in some cases the process takes 50 negotiation steps (as stated before, we report results here do refer to the monotonic, time-dependent concession case). Thus, overall these results support the underlying idea put forward in this chapter. That is, having a maximal utility graph of possible interdependencies can be used to successfully navigate the contract space and reach Pareto-efficiency with a limited number of steps, even for a relative large number of interdependent issues.

Finally, we should mention that, while our method for generating random graphs is well grounded in theory and allows us to generate highly non-linear preference classes, future work we could consider adapting custom-made generators, such as CATS (Leyton-Brown et al. [142]) for this problem. While the CATS approach was designed mostly for generating bids in combinatorial auctions, it could conceivably be adapted for multi-issue bargaining settings, such as ours.

3.6.4 Experimental set-up and analysis of the collaborative filtering model

All the experimental results reported above for automated negotiation tests refer to the case when the seller does know an approximate (i.e. maximal) structure of the utility graph of the buyer (although, of course, not anything about the actual values in that graph, for the particular buyer he is negotiating with). Otherwise stated: even though for each negotiation trace both the structure and utility values for each buyer were generated at random, the seller agent does know a graph topology which is a super-graph of the actual topology of the buyer (with a certain number of extra, spurious edges added at random). The question is: in a real e-commerce scenario, how does the seller acquire this (approximate) structure information in order to help him make proposals during the negotiation? The answer (already extensively provided in Section 3.5 of this chapter) involves using the intuition that graph topologies for a set (or class) of buyers are not completely random, but largely overlapping. Therefore, by looking at previous concluded negotiation data, the seller can approximately retrieve this structure. In this section, we present the results from experiments performed to test this graph topology retrieval model.

The settings used in generating artificial negotiation data for graph retrieval are largely the same as the ones described in the above section. Thus, we still work with graphs containing 50 issues, 80 random binary dependencies of which 50 positive and 30 negative. The tests we performed can be basically classified into two parts, which follow the structure of the model presented in Section 3.5:

- First, a set of tests was performed in order to choose which similarity criteria are relevant and useful for this setting (see Sect. 3.5.5). In this section, we examined how robust are these results to higher deviations in generating random buyer profiles and how much previous negotiation data is needed in order to efficiently learn this approximate structure.
- The second set of tests concerned the selection of the cut-off point for the number of edges to be included in the maximal graph (see Sect. 3.5.7 for a formal description). We tested different graph sizes and we examined how error in the retrieval of the structure of the maximal graph affects the negotiation stage.

In the following, we first describe the experimental set-up for these tests. Next, in subsection 3.6.5, we present and compare the results for the two collaborative filtering criteria considered. Finally, in Section 3.6.6 we show the results for determining the maximal graph cut-off point.

Experimental set-up for graph retrieval tests

There are two dimensions across which the efficiency of retrieval needs to be tested:

- **The strength of the interdependencies in the generated buyer profiles.** This is measured as a ratio of the average strength of the inter-dependency over the average utilities of an individual item. To explain, each buyer profile is generated as follows:
First, for each item, an individual value is generated by drawing from identical, independent normal distributions (i.i.d.-s) of center $C_{individual-item} = 1$ and variance 0.5. Next, the substitutability/complementarity effects for each binary cluster are generated by drawing from a normal i.i.d.-s with a centers $C_{non-linearity}$. The strength of the interdependency is then taken to be $\frac{C_{non-linearity}}{C_{individual-item}}$. The smaller this ratio is, the more difficult it will be to detect non-linearity (i.e. complementarity and substitutability effects between items). In fact, if this ratio takes the value 0, there are no effects to detect (which explains the performance at this point), at 0.1 the effects are very weak, but they become stronger as it approaches 1 and 2 (which is the range considered here). Thus, this ratio can be intuitively interpreted as a “signal to noise” ratio (i.e. the non-linearity effect representing the signal to be detected, vs. the noise due to dispersions in individual utilities).

- **Number of previous negotiations for which information is available.**

The performance measure used is computed as follows. Each run of the collaborative filtering algorithm (for a given history of negotiations, and a certain probability distribution for generating that history) returns an estimation of the utility graph of the buyer. Our performance measure is the recall, i.e. the percentage of the dependencies from the underlying utility graph from which buyer profiles are generated, which are found in the graph retrieved

by the seller. Due to noise and/or insufficient data, we cannot expect this graph retrieval process to always have 100% accuracy. We therefore studied what is the effect of an imprecise graph on the part of the seller on the negotiation process itself (stage 2 of our approach). This is discussed in Section 3.6.6.

3.6.5 Results from retrieval experiments using cosine-based vs. correlation based similarity

The setting presented above was tested for both cosine-based and correlation based similarity (see Sect. 3.5.5). Figure 3.8 gives the resulting graphs for the cosine-based case, while Fig. 3.9 gives the results for the correlation-based one. Each of the points plotted and resulting dispersions was computed by averaging over 100 different tests. To make the tests as independent as possible, a new data set was generated for each test.

The overall conclusion which can be drawn from our tests (see Fig. 3.8 and Fig. 3.9) is that one of the techniques we investigated, namely correlation-based similarity is considerably more successively than the simpler, cosine-based similarity technique. This can be observed from Fig. 3.8 and Fig. 3.9: while correlation-based similarity can extract 96% (+/- 7%) of dependencies correctly given enough data (from around 1500 completed negotiations) and strong enough dependency effects (above 1), cosine-based similarity achieves a maximum of just above 40%. Thus, we find that correlation-based similarity seems to perform well for our task - i.e. retrieving the graph topology, especially since we need to detect not only complementarity effects, but also substitutability ones. Cosine-based similarity (which, according to some sources, e.g. [145] was the initial algorithm behind Amazon's recommendations) is conceptually simpler and works well only in detecting complementarity dependencies and only in the case when the data is relatively sparse (each buyer expresses interest only in a few items). Correlation-based similarity gives a more sophisticated measure which does not have these limitations.




Figure 3.8: Results for the cosine-based similarity. Left-side graph gives the percentage of correctly retrieved dependencies, with respect to the average interdependency strength, while right-side graph gives the percentage of correctly retrieved dependencies with respect to the size of the available dataset of past negotiation traces.




Figure 3.9: Results for the correlation-based similarity. Left-side graph gives the percentage of correctly retrieved dependencies, with respect to the average interdependency strength, while right-side graph gives the percentage of correctly retrieved dependencies with respect to the size of the available dataset of past negotiation traces.

3.6.6 Experimental results for selecting graph cut-off number of edges in the maximal graph

After measuring the effect of the two similarity criteria considered (i.e. cosine and correlation-based), as well as the effect of different amounts of data, we present results for different cut-off sizes for the maximal graph (i.e. the k parameter introduced in Section 3.5.7). For all tests reported in this Section, we used correlation-based similarity and we assumed 1000 records of previous negotiations are available for filtering. We chose to focus on correlation-based similarity since this criteria clearly performs better, in this setting, than cosine-based similarity (see above results). Also, as shown in Sec. 3.6.5, 1000 records is a reasonable amount of data to ensure a good accuracy of retrieval for correlation similarity.

For all the tests reported here, we report the cut-off values (which are in fact, a maximal number of edges considered) as percentages of the number of edges in the true, underlying graph of the buyer (which, as shown above, contains 80 edges, generated at random).




Figure 3.10: Percentage of correctly retrieved dependencies from the underlying graph of the buyer for different number of cut-off number of edges considered. On the vertical axis, the difference to 100% corresponds to the percentage of missing edges in the retrieval. The number of cutoff edges on the horizontal axis are given as percentages of the actual true size of the buyer's graph (i.e. 80 edges in our case)

From Figure 3.10 we can see that the number of missed edges decreases as we increase the number of edges taken as part of the maximal graph (the edges are taken in decreasing order of value from the correlation tables). However, we should point out that after a value of 150% - 200% from the actual size of the true graph (i.e. $k = 3$ from Section 3.5.7), this increase is not so great and the dispersion of the results also increases. Intuitively, this means that there are a number of edges - about 15%-20% of the total (remember graphs are generated at random), which appear inherently “hard” to find for the filtering algorithm.

Of course, we may achieve a higher percentage if we take more information on concluded negotiations, but for consistency, here in all tests we limit ourselves to 1000 records.

After determining what the error rate is likely to be for different maximal graph sizes (i.e. sizes of k), we performed another set of tests in order to estimate the loss in Gains from Trade during the actual negotiations when the seller starts the negotiation with a graph topology which misses edges or contains considerably more edges than usual. Results are reported in Figures 3.11 and 3.12. For all results reported in these figures, 100 tests/point have been performed. We should point out, though, that the difficulty of the search problem in such a setting depends not only on the sparseness or density of the graph, but also on the dispersions of the normal utility function used to generate random values in the clusters corresponding to each edge (as discussed in the experimental results from the negotiation model itself). In order not to artificially inflate the results, for all the tests reported here we used a normal distribution centered around 1 with spread 5, which is the most non-linear case considered during the negotiation tests.

From Figs. 3.11 and 3.12, several conclusions can be drawn. First, missing edges from the graph the Seller starts the negotiation with has a considerably greater negative effect than adding too many extra (erroneous) edges.

Thus, as shown in Fig. 3.11, in order to get above 90% of the optimal Gains from Trade in future negotiations, the retrieval process cannot miss more than about 15% of the true inter-dependencies in the true graph of the Buyer. However, having a considerably denser starting graph does not degrade the performance so significantly. In fact, as we see in Fig. 3.12, having 3 times as many edges than in the original buyer graph (which means 2/3 of all edges are erroneous), only decreases performance with around 4%.

By examining the 3 graphs above, we can conclude that in this setting (i.e. a random graph of 50 issues with 80 edges), the best cut-off point would be having 200% more dependencies than in the true graph of the seller (i.e. around 150 edges, or a $k = 3$). This level would mean that about 15% from the edges in the true super-graph of the buyer will be missed by the filtering process. However, it would ensure that we still get at least 90% of optimal Pareto-efficiency, on average, after 40 negotiation steps.

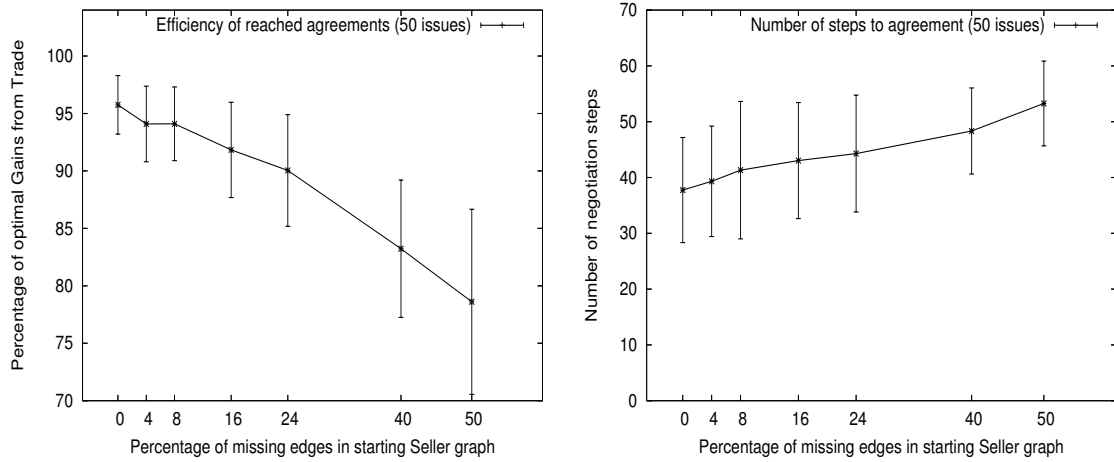


Figure 3.11: Effect of missing edges (dependencies) in the starting Seller graph on the Pareto-optimality of reached negotiation outcomes




Figure 3.12: Effect of excess (erroneous) edges in the starting Seller graph on the Pareto-optimality of reached negotiation outcomes

3.7 Discussion

In this section we provide a comprehensive review of related work, highlighting the features which are comparable to those from our approach. First, we position our work on multi-issue negotiation, by comparing it to other direct-revelation allocation methods such as combinatorial auctions. We show there is a strong connection between our approach with the work on preference elicitation over multiple issues (items). Next, we provide a review of previous work on agent-mediated negotiation, focusing in particular on multi-issue negotiation models (Sect. 3.7.1). We conclude by summarizing the main contributions of our work and identifying directions for future research.

3.7.1 Comparison to other automated negotiation (bargaining) approaches

Agent-mediated negotiation is a widely researched problem in electronic commerce and resource allocation settings (for comprehensive overviews of the field, the reader is asked to consult, e.g.: Kraus'99 [129], Lomuscio et al.'03 [147], Gerding et. al, '00 [82], Lai et. al, '05 [134]). Several previous results model automated negotiation as a tool for supporting the buyer's decision process in complex e-commerce domains [60, 79, 126, 207]. Most of the work in multi-issue negotiations has focused on the independent valuations case.

Faratin, Sierra & Jennings [71] introduce a method to search the utility space over multiple attributes, which uses fuzzy similarity criteria between attribute value labels as prior information. Coehoorn and Jennings [163] extend this model with a method to learn the preference weights that the opponent assigns to different issues in the negotiation set, by using kernel density estimation. These papers have the advantage that they allow flexibility in modeling and deal with incomplete preference information supplied by the negotiation partner. They do not consider the question of functional interdependencies between issues, however.

Other approaches to multi-issue negotiation problem are the agenda based approach (Fatima et. al. [72, 73]) and the fuzzy constraint-based negotiation approach (Luo et. al. [149, 232]). Debenham [60] proposes a multi-issue bargaining strategy that models the iterative information gathering which takes place during the negotiation. The agents in Debenham '04 [60] do not explicitly model the preferences of their opponent, but construct a probability distribution over all possible outcomes. However, these models are not explicitly designed to address the problem of complex and high dimensional negotiations.

The argumentation approach to negotiation (see Rahwan et. al. [178] for an overview) allows the agents to exchange not only bids, but also arguments that influence other agents' beliefs and goals, which, it is claimed, allows more flexibility. Some issues which are usually left open in such approaches are: how do the agents' mental states relate to their utilities and if (or how) can the efficiency of such negotiations be measured from game-theoretic perspective.

Raz et. al.'06 [183] propose a model for bilateral multi-issue negotiation with bounded

rational agents (including humans). Their approach uses Bayesian updating to estimate the type of the opponent, considering linearly additive preference functions. Kraus & Schechter, '03 [128] consider a model of bilateral negotiation, but in a complete information setting. Their work mostly focuses on examining agent strategies in case one agent loses utility over time, while the other gains utility over time.

Luo et. al. [148] propose a “default and adjust” technique to acquire a user’s trade-off preferences for negotiating agents. This work is only loosely related to ours, since the paper mostly focuses on acquiring trade-offs between two, continuous-valued issues (e.g. quantity and price). Gerding et. al. [79] also consider two attributes (price & quantity) in a setting where agents utilities can be described by a Cobb-Douglas utility function. The use evolutionary simulation to model the agents’ learning and adaptation. The emphasis of this paper is more on adaptation to different strategies in a one-many setting than on multi-issue exploration. Other lines work focus on other aspects of negotiation, such as coordination between multiple, concurrent negotiation threads (Nguyen & Jennings [167]), generating promises of future rewards in repeated negotiation games (Ramchurn et al. [181]) or possibilistic selection of negotiation partners (Brzostowski & Kowalczyk [34]). Other lines of work propose multi-issue negotiation protocols specifically geared to an application domains, such as crisis management (Hemaissia et al. [99]) or internet-based service provision (Dang & Huhns [57]).

Somefun, Gerding & La Poutré [80] propose a negotiation mechanism where the bargaining strategy is decomposed into a concession strategy and a Pareto-search strategy. This is somewhat similar to the approach taken in this chapter, though here we mostly focus on the Pareto-search. The Pareto search strategy discussed in [80] is, however, rather different than the one presented on this chapter, since [80] focuses on two continuous issues, corresponding to a two-part tariff. The search is performed using an orthogonal or derivative-follower strategy, in the space of ISO-utility lines. A similar negotiation strategy, which involves choosing an offer on the iso-utility (or “indifference”) curve that is closest to the best offer made by the opponent in the previous negotiation period is proposed by Lai, Sycara & Li [135, 138]. The methods presented in these papers build on earlier work by Ehtamo & Hamalainen [67], who present a constraint proposal method to construct the Pareto-efficient frontier of a multi-attribute negotiation, by means of a non-biased mediator.

Boutilier et al. [26] present a cooperative negotiation model for autonomous systems, through incremental utility elicitation. Using decentralized resource allocation as a problem setting, they emphasize the difficulty of eliciting complex utility functions and propose a strategy that requires only a small set of sampled utility function points in order to find near-optimal allocations. The perspective on negotiation as partially a preference modeling problem is close to the one taken in this chapter - however there are important differences, since their model uses minimax regret, not Pareto-efficiency as an optimality criteria and the work reported in [26] does not handle multidimensional utility functions over multiple resources.

Another, more fundamental, direction is taken in papers that study the communication complexity (e.g. Chevaletre et. al. '05 [47, 48]) and or computational complexity (Dunne et.

al. [66]) of multi-issue negotiation protocols. The work of Chevalleyre et al. [47, 48] is related to our approach in the way they define complex utility functions over bundles of resources. However, unlike this work (and much of the work on multi-issue negotiation reviewed in this Section), their work is mostly concerned with deriving theoretical bounds for different negotiation protocols and utility function classes. While this provides very useful insight for other research, they do not propose directly implementable negotiation algorithms or heuristics (that can handle incomplete information settings, for instance).

Two negotiation approaches that specifically address the problem of complex inter-dependencies in high dimensional negotiations over binary issues — and are therefore most related to our work — are Klein et. al. [126] and Lin [144]. Klein et. al. [126] use a setting similar to the one considered in this chapter, namely bilateral negotiations over a large number of boolean-valued issues with binary interdependencies. In this setting, they compare the performance of two search approaches: hill-climbing and simulated annealing and show that if both parties agree to use simulated annealing, then Pareto-efficient outcomes can be reached. In a similar line of work, Lin [144] uses evolutionary search techniques and mediation to reach optimal solutions. By comparison to our work, these approaches do not try to use prior information, in the form of the clustering effect between the preference functions of different buyers, in order to shorten individual negotiation threads. At least from the results reported in [126], they appear computationally more expensive (around 2500 steps / 50 binary issues). However, there is no established benchmark for this problem, so the complexity of non-linear dependencies considered in [126] cannot be directly compared to the complexity considered here (though, in future work, it would be interesting to compare the two approaches, if a clear benchmark could be defined).

In more recent work, Ito, Hattori & Klein [108] propose an auction-based negotiation protocol for agents with non-linear utility functions. In their approach, each agent samples and searches its own utility space, and then submits the best found contract points to a central mediator (or auctioneer). The mediator then selects the contract combination that is consistent and maximizes the social welfare of the bidders. Although such a centralized method (resembling a combinatorial auction) has some advantages (e.g. it easily allows multi-party negotiation), assuming a central, impartial mediator may not be suitable in many application settings. Fujita & Ito [75] develop an extension of this work, by proposing a mechanism of adjusting the utility threshold above which bids are accepted by the auctioneer.

Another related work, which also considers bundles of binary issues is Somefun, Klos and La Poutré [207]. Our work builds on the framework for analyzing negotiation outcomes developed in [207], especially with regards to defining Pareto-utility through Gains of Trade (see Appendix 3.A of this chapter). However, the actual negotiation model proposed in [207] involves attempting to learn a relationship between pairs of *bundles* of items (as opposed to between the items themselves, such as in this work). As the number of bundles is exponential, even for relatively small number of items (e.g. for 10 items, there are already $2^{10} > 10^6$ possible bundles), their approach is computationally feasible only for small-scale settings.

Hyndriks, Thykohonov & Jonker [102] handle bilateral multi-issue negotiations with interdependent valuations, by proposing a method to eliminate the dependencies between

the issues. Unlike our model, their approach does not involve building a model of opponent preferences, but it relies on approximating a jointly agreeable point in the contract space and then performing a local search around this point (through a weight function). This may be a promising solution in some settings, although in the complex k -additive utilities case we consider, eliminating or “linearising” the dependencies would not be feasible for most settings. Furthermore, in a large contract space finding a good initial contract (around which to base the search) is a complex problem in itself, which, similar to our work, requires some implicit assumption about the structure of buyer preferences.

3.7.2 Relation to graphical utility models and preference elicitation

Regarding graphical representation of utilities, several other well-known formalisms have been proposed and analysed in terms of expressiveness and conciseness of the functions which can be represented. The best known are CP-nets and Directed Acyclic Graphs (DAGs) - and their extensions [9, 30, 32, 39, 44, 219]. Some of this work was actually developed in parallel to ours, with a view of application to a different class of problems although, in retrospect, it would be very interesting to also study how these formalisms could be used in a multi-issue negotiation setting. There are, however, some important differences by comparison to our utility graphs definition.

First, existing literature on DAGs uses *directed graphs* to model utility. By comparison to [9, 30, 39, 44], our model is based on undirected graphs, which are specifically geared for representing k -additive utility functions. In the setting of negotiation, the undirected utility graph formalism may have some important computational benefits. As shown in this chapter and the related publications, efficient heuristics can be found (c.f. Sect. 3.4) for decomposing such utility graphs online and learning the associated cluster values. Moreover, for the undirected graphical model proposed here, we also show how the structure of the graph itself could be constructed from past, aggregate negotiation data. To the best of our knowledge, this has not been attempted before in prior literature on directed graphical utility models. In Bayesian probability theory, learning causalities from data is known to be a considerably more difficult than just learning correlations [172, 173]. Therefore, we can conjecture that learning structure could also prove more difficult for directed utility graphs than for undirected ones, although further research is needed in this direction.

Another important and somewhat related line of work is that of preference elicitation. In this context, we can identify a link between models used in preference elicitation and multi-item negotiation. For example, Brazunias and Boutilier [32] propose a model (developed independently and concurrently with our work), for utility elicitation in generalized additive independence (GAI) models. There are, however, also important differences. First, the model proposed in [32] uses a different graphical formalism to encode preferences, while their optimization criteria is the error in reached in the utility of the buyer, not Pareto efficiency. In preference elicitation settings, the prices asked by the seller are fixed throughout the process and are known to the oracle (if one sees the oracle as roughly corresponding to the buyer agent in our setting, i.e. the party whose preferences need to be modeled or

learned). By contrast, negotiation is usually a game with double-sided incomplete information. Although the parties start from an initial vector of asking prices, the buyer has an incentive not only to explore the bundle contents, but also to bargain about the price of the bundle being offered (according to his/her negotiation strategy). Finally, in [32] (just as in our previous work [186]), the problem of acquiring the initial graphical model structure is left for future research. In this chapter a complete model is proposed, since we show how the topology of such utility graphs (an implicitly of the utility space) can be approximated from anonymous buyer data, using techniques from item-based collaborative filtering [145].

3.8 Conclusions and future work

Much of the existing work (of which we are aware) of that consider negotiation settings close to ours (i.e. large number of issues with complex dependencies [108, 126, 144]) assume, in some form, the presence of an independent mediator. The role of the mediator and the exact protocol used varies in each approach, corresponding to varying degrees of centralization. In settings where a non-biased mediator is used (similar to combinatorial auctions), the problem of indirectly learning the negotiation opponent's preferences is not posed (and often does not need to be posed), because the mediator agent can elicit the preferences of negotiating agents through direct queries. This is a significant difference with this work, where the problem of modeling a randomly-encountered negotiation partner's preferences is central in order to reach jointly efficient (close to Pareto-optimal) agreements. Given the complexity of this task, we propose a solution based on two-step model. First, we use some information about the most likely preference structure, based on the clustering effect of the preferences for a statistical population of buyers (a similar approach as to that taken in collaborative filtering). Next, this model is refined and improved during the negotiation stage itself, based on the counter-proposals of the specific negotiation partner.

The contribution of our work to existing negotiation literature can be summarized as follows:

- It shows that multi-issue negotiations for agents with complex, k -additive utility functions can lead to jointly efficient outcomes, even with a limited number of negotiation steps are available (due to time constraints, buyer impatience etc.) and even when the preference functions of the two parties remain private.
- It considerably improves the speed of existing heuristics for this problem, for the type of utility functions we consider, by using anonymous, aggregate buyer data.
- It shows that graphical models of utility (an area of AI that has received considerable research attention) can be naturally applied to multi-issue negotiation settings. In particular, it shows the usefulness of undirected graphical models in handling negotiations for agents with k -additive utility functions.

- It establishes a formal link (in the form of the maximal structure of a utility graph for a class of buyers) between opponent preference modeling used in negotiation settings and the way of modeling preferences used in collaborative/social filtering systems.

As future work, there are several directions which could be explored in this area. An immediate one is to study other classes of non-linear preference functions for which it is possible to reach Pareto-efficient agreements, under double-sided incomplete information, with a linear number of negotiation steps. To this end, we intend to make use of results from random graph theory [191] and constraint processing [61].

Second, we could consider several, distinct super-graphs for different sub-populations of buyers (rather than just one, as in this chapter). Buyers could then be assigned to a certain sub-population at runtime, during the negotiation thread itself. In the longer term, another potentially very fruitful area of research would be to explore the connection between our work and problems studied in preference elicitation. Arguably, the techniques developed in this chapter and [186] in the context of multi-issue negotiation could also be applied to the problem of eliciting user preferences for non-linear, high-dimensional settings.

Appendix 3.A: Equivalence between maximization of the Gains for Trade and Pareto-optimality

In this Appendix, we formally prove that, in quasi-linear utility settings, bundles maximizing the Gains from Trade Pareto-dominate all other bundles. This proof holds for settings where the utilities of the agents can be expressed in monetary terms, i.e. are *quasi-linear*. For our case, this means the utility that buyer and seller get from trading a bundle of items can be expressed in terms of *net monetary value*. For the buyer, this net monetary value is the difference between his utility from getting the bundle being exchanged (expressed in terms of money) and the price he pays for it. For the seller's (i.e. electronic merchant's), the net monetary value is the difference between the price he receives for a bundle and his cost for providing it. This proof first appeared in previous work on negotiation of [207].

Formal Discussion

Before being able to more formally state the results, some notation is necessary. Let $N \subset \mathbb{N}$, with $n = |N|$, denote the collection of n individual goods and 2^N denotes the power set of N (i.e., the collection of all subsets of n), then $B = 2^N \setminus \{\emptyset\}$ denotes the collection of all possible bundles. Furthermore, let $P = \mathbb{R}$ denote the collection of all possible bundle prices.⁴ The customer and the shop attach the monetary values of $v_c(b)$ and $v_s(b)$, respectively, to a bundle $b \in B$ (with $v_c(b), v_s(b) \in P$). The function $x_j : B \times P \mapsto \mathbb{R}$ with $j \in \{c, s\}$ denotes the *net* monetary value for bundle b and bundle price p : $x_c(b, p) = v_c(b) - p$ and $x_s(b, p) = p - v_s(b)$ denote the customer's and the shop's net monetary values, respectively.

⁴Negative prices may not be realistic, but we want to make as few behavioral assumptions as possible. For the results the possibility of negative prices is not problematic (see Footnote 5).

We assume that the customer's and the shop's utility for consuming bundle b for a price p , denoted by $u_j(b, p)$ with $j \in \{c, s\}$, can be expressed as the composition function $g_j \circ x_j(b, p)$ with $j \in \{c, s\}$ and $g_j : \mathbb{R} \mapsto \mathbb{R}$. For g_j we assume that $\frac{dg_j(x)}{dx} > 0$ for all $x \in \mathbb{R}$ and $j \in \{c, s\}$. Thus we have that $u_j(b, p) = g_j(x_j(b, p))$ and since g_j is a strictly increasing function we can without loss of generality assume that $u_j(b, p) = x_j(b, p)$ for $j \in \{c, s\}$ (cf. [153]).

Given the customer's and shop's monetary values, we define a useful subset B^* of B as follows: $B^* \equiv \arg \max_{b \in B} (v_c(b) - v_s(b))$, that is, B^* represents the collection of bundles with the highest gains from trade. We are now ready to introduce the following proposition.

Proposition 2 *A deal (b, p) with $b \in B$ and $p \in P$ is Pareto efficient if and only if $b \in B^*$.*

Remark 1 *A deal (b, p) is Pareto efficient if there is no (b', p') such that $u_j(b, p) \leq u_j(b', p')$ for all $j \in \{c, s\}$ and the inequality is strict for at least one j .*

Proposition 2 means that a deal is Pareto efficient if and only if it entails a bundle with the highest gains from trade. For the proof of this proposition the following lemma is very useful.

Lemma 1 *For any two deals (b^*, p^*) and (b, p) with $p^*, p \in P$, $b^* \in B^*$, and $b \in B \setminus B^*$ we have $x_c(b, p) < x_c(b^*, p^*)$ or $x_s(b, p) < x_s(b^*, p^*)$.*

Proof 1 *We prove the above lemma by contradiction. Suppose that for any $b^* \in B^*$ and $b \in B \setminus B^*$ we have $x_c(b, p) \geq x_c(b^*, p^*)$ and $x_s(b, p) \geq x_s(b^*, p^*)$. A necessary conditions for this to hold is that $v_c(b) - v_s(b) \geq v_c(b^*) - v_s(b^*)$. However, $b^* \in B^*$ and $b \in B \setminus B^*$ means, by definition of B^* , that $v_c(b) - v_s(b) < v_c(b^*) - v_s(b^*)$.*

We are now ready to prove Proposition 2.

Proof 2 1. (If) Pick any $j \in \{c, s\}$. Suppose that j 's position improves by moving from any deal (b, p) with $b \in B^*$ to (b', p') , that is, $u_j(b, p) < u_j(b', p')$. It then suffices to show that the opponent denoted by j' will always be made worse off, that is, $u_{j'}(b, p) > u_{j'}(b', p')$. From the properties of g_j and $g_{j'}$ it follows that a bargainer's position improves/worsens whenever the net monetary value increases/decreases. Since j 's position improves, it follows from Lemma 1 that j' is made worse off whenever $b \in B \setminus B^*$. Moreover, if $b^*, b \in B^*$ then the gains from trade remain unchanged, hence j' is made worse off.

2. (Only if) We will prove this part by contradiction. Suppose that $b \notin B^*$ with the price being any $p \in P$. Pick any $b' \in B^*$ and set the bundle price to $p' = p + v_s(b') - v_s(b)$, so that $p' - v_s(b') = p - v_s(b)$. It follows from $p \in P$ that $p' \in P$ (recall that $P = \mathbb{R}$)⁵

⁵If we choose to a priori rule out $p < 0$ and $v_j(b) < 0$ (for $j \in \{c, s\}$ and all $b \in B$), then $p \geq v_s(b)$ should hold because otherwise the shop will not be willing to sell the bundle in the first place. Consequently, $p' \in P$ still holds.

and the properties of g_s that the shop is indifferent between the deals (b, p) and (b', p') . Also, it follows from Lemma 1 and the properties of g_c that the customer is made better off. That is, any $b' \in B^$ Pareto dominates $b \notin B^*$. Thus $b \notin B^*$ cannot be a Pareto efficient solution.*

Part II

Preferences under uncertainty and strategic reasoning in sequential auctions

Chapter 4

Designing bidding strategies in sequential auctions for risk averse agents

4.1 Introduction

Design of electronic auctions is considered an important open area of research in electronic commerce, both from a theoretical and an application perspective. There are two main approaches to this problem. One concerns the design of the auction mechanism itself, such as it guarantees certain properties, such as efficiency, individual rationality or budget balance. However, for some auction designs, such as simultaneous ascending, sequential and repeated auctions, this is not possible and research has focused on designing the bidding strategies of the agents participating in such auctions.

As previous shown in [27, 89, 184, 217], the main problem that a bidder has to face in a sequential (or simultaneous ascending) auction is the exposure problem. Informally stated, exposure means that an agent has to commit to buying an item (and thus take a “sunk” cost [184]), before she can be sure that she will be able to secure other items in her useful set or bundle (i.e. the set of items that gives her a positive utility). If she does not manage to acquire the other items, she is exposed to the risk of a loss.

In order to deal with this problem, several strategies have been proposed in existing literature. Boutilier et al. '99 examines the role of dynamic programming in computing bidding policies in sequential auctions, based on distributions over estimated prices. Reeves et al. '03 [184] study the problem of bidding in simultaneous ascending auctions (a problem closely related to the sequential settings) - in the context of market-based scheduling. Osepshvili et al. '05 [1] continue this line of research, but use probabilistic prediction methods

of final prices and introduce the concept of self-confirming price distribution predictions. Gerding et al.'07 [83] derive the optimal bidding strategy for a global bidding agent that participates in multiple, simultaneous second-price auctions with perfect substitutes. Unlike this work, however, they do not consider complementarities (i.e. agents requiring bundles of items), and the setting is slightly different, as all auctions are assumed to close exactly at the same time, not sequentially.

In a direction of work that considers a setting very related to this chapter, Greenwald & Boyan '04 [89] study the bidding problem, both in the context of sequential and simultaneously ascending auctions. For the sequential auctions case, they consider a decision-theoretic model and show that marginal utility bidding represents an optimal policy. Their result applies, however, only to risk neutral agents. Hoen et al. '05 [217] look at the related problem of bidding in repeated auctions with complementarities and draw a parallel with the N-person iterated prisoner's dilemma. The above approaches have been shown to be efficient in many situations, both in self play and against a wide variety of other strategies, in competitions such as the TAC. Although most do implicitly consider the aspect of risk, they do not explicitly model the risk-taking *attitude* of the bidding agents. By "explicitly model" we mean building a profile of the agent's risk preferences towards uncertain, future outcomes (such as the final allocation of a sequential auction).

In standard economic theory, since the seminal work of K. Arrow and J. Pratt, preferences towards risk have been considered essential in understanding and modeling decision making under uncertainty [5, 88, 122, 171]. In fact, a body of auction theory from economics [158, 171] identifies risk preferences as a very important, open research area. In recent econometrics and financial economics literature, this has led to considerable research interest in efficiently modeling and eliciting risk aversion from human users [45, 171, 176]¹ Existing economic approaches to risk modeling do not, however, consider sequential auctions over combinations of items, nor propose bidding heuristics for this setting.

From the point of view of multi-agent systems literature, only a limited number of papers discuss risk profiles. Babanov et al. '04 [8] use the concept of certainty equivalence, similar to our work, in the context of optimal construction of schedules for task execution. Liu et al. [146] do consider risk-aversion on the part of the agents (similar to the approach taken in this chapter) - but their work is mostly concerned with providing an analytical solution to the one-shot auction case. Vytelingum et al '04 [58] consider risk-based bidding strategies in a double-auction setting. However, both the auction setting (i.e. CDA) and the risk model used (which is not based the standard Arrow-Pratt model) make this work rather different in focus from ours. Finally, Vetsikas and Jennings [223, 224] also consider a model that includes agent attitudes towards risk (among other factors, such as budget constraints and reserve prices), for the case on multi-unit, sealed-bid auctions. They provide a thorough theoretical analysis of this case, but they do not consider complementarities (i.e. agents desiring bundles of goods), nor sequential allocation.

¹A practical example of risk elicitation in finance are the questionnaires involving probabilistic choices between several scenarios that investment fund managers send to potential investors.

4.1.1 Goals and organisation of this chapter

The basic goal of this chapter is to study the relationship between a bidder agent's attitude towards risk (measured by the standard Arrow-Pratt risk aversion model - more specifically the CARA model) and her perceived best available bidding policy in a sequential auction (modeled by a Markov Decision Process). In this context, we consider the bidding decisions of an agent that desires a bundle of complementary-valued goods that are sold through a sequence of auctions.

First, we investigate analytically how an agent's perception of her optimal bidding policy, given her probabilistic expectation of future prices, is affected by her risk aversion profile. Similar to [27, 89, 112, 184, 217], we take a decision-theoretic approach to the design of bidding agents, meaning agents reason w.r.t. the probability of future price distributions, and do not explicitly deliberate over the preferences, risk profiles and strategies of other bidders. Next, we conduct an experimental study of how an agent's with complementarities attitude towards risk affects her chances of winning a desired bundle, when bidding against a population of local bidders, desiring only one good. Furthermore, we also look at how this bidding policy affects the auctioneer's revenue. Our primary goal is to gain a qualitative understanding of how sequential auction markets are influenced when bidders with complementary valuations participating in them are risk averse.

The remainder of the chapter is organized as follows. Section 4.2 presents the risk aversion model, which forms the foundation of the following sections. Section 4.3 describes the bidding model and discusses the optimal bidding policies for both first and second-price sequential auctions. Section 4.4 provides the experimental results, while Section 4.5 concludes the chapter with a discussion.

4.2 Modeling Utility Functions Under Risk

The literature on risk aversion identifies several 3 main types of agents w.r.t. their risk profiles: risk averse, risk neutral (indifferent) and risk proclive ("risk loving") agents. In the following we will focus our attention mostly on the risk averse and risk neutral cases, since these are the cases that describe the behaviour of economic agents in most practical situations (c.f. [5, 158, 171]). Denote the private payoff z achieved by an agent participating in an auction or lottery. The utility a risk-averse agent assigns to this payoff is described by the Arrow-Pratt utility function:

$$u(z) = 1 - e^{-rz} \text{ for } r > 0 \quad (4.1)$$

For the case of risk indifference ($r = 0$), we take $u(z) = z$.

Note that the auction model we consider in this paper is a *private value* model. The payoff z of a bidder after participating in an auction is a difference between a private value v and the amount of money paid to acquire the item in the auction (or cost) c . Since the private

value v is private to each agent, the payoff value is also private. Therefore, in a risk averse setting there are parameters describing the private preferences of a bidder: value v and risk aversion coefficient r . Preferences of agents cannot be directly compared by comparing their private values or payoffs, as the risk factors r must also be taken into account.

Our choice of defining Eq. 4.1 represents a standard form of defining utility functions under uncertainty [171] (the same choice is made in [159, 171], among others). This form ensures that the following relation holds:

$$r_u(z) = -\frac{u''(z)}{u'(z)} \quad (4.2)$$

As defined in Eq. 4.2, $r_u(z)$ corresponds to the Arrow-Pratt measure of absolute risk aversion [5, 171]. In this chapter, we consider r constant for each agent, i.e. $r_u(z) = r, \forall z$, thus we use the constant absolute risk aversion (CARA) model.² Factor r represents a constant which differs for each agent, characterizing her own preference towards risk-taking.

We use a state-based representation, in which all possible future outcomes at time t is denoted by S_t . All $s \in S_t$ are assigned by the agent a monetary payoff z_s and an expected probability p_s (where $p_s > 0$ and $\sum_{s \in S_t} p_s = 1$). We define the lottery L_t over a set of payoffs \tilde{z}_s (corresponding to the state \tilde{S}_t) as the set of payoff-probability pairs, i.e. $L_t = \{(z_s, p_s)\}$ where $s \in S_t$. In this form, the definition is generic, but as we show in Sect. 3, there is a natural correspondence between lotteries and states in a sequential-auction game.

The *expected utility* of the agent at time t over the lottery L_t is described by a von Neumann-Morgenstern utility function:

$$E_u[L_t] = \sum_{(z_i, p_i) \in L_t} p_i u(z_i) \quad (4.3)$$

In case all the agents are risk-neutral (i.e. have $u(z) = z$), it is easy to compare expected utilities and payoffs across agents. However, for risk averse agents this is not the case, and we need a measure that enables comparison of payoffs across agents with different attitudes to risk in uncertain domains. The utility functions of the agents are not directly comparable in this setting, since each agent has a different attitude towards future risk (different r factor).

The widely used concept in risk modeling is to identify a monetary value (i.e. amount of money), such that the agent is indifferent between receiving this value with certainty or entering the lottery. This amount is called the **certainty equivalent (CE)** of the lottery. It can be seen as the monetary payoff the agent would attach to the future, if all the uncertainty (and hence risk) were discounted.

Formally defined, the **certainty equivalent (CE)** of a lottery L_t is defined as the certain payoff value which is equivalent to the expected utility of the lottery L_t . That is:

$$u(CE(L_t)) = E_u(L_t)$$

²This is a widely used risk aversion model, which we deemed sufficient for the purpose of this work. We leave the study of Relative Risk Aversion (RRA) models to future research.

Expanding both sides using Eqs. 4.1 and 4.3 above, we have:

$$-e^{-rCE(L_t)} = \sum_{(z_i, p_i) \in L_t} -p_i e^{-rz_i}$$

Hence the following expression can be derived for the certainty equivalent of the lottery:

$$CE(L_t) = \begin{cases} -\frac{1}{r} \ln \sum_{(z_i, p_i) \in L_t} p_i e^{-rz_i} & \text{for } r > 0 \\ \sum_{(z_i, p_i) \in L_t} p_i z_i & \text{for } r = 0 \end{cases} \quad (4.4)$$

In other words, the certainty equivalent can be seen as the certain amount of money which has the same utility to the agent as the equivalent lottery, before the outcome of the lottery is known. In the following, we define and prove a recursive property of CE functions, which is relevant for their application to sequential games considered in this paper.

Property 1: Suppose we have a game that occurs in stages t ; at each time step t the game can transition into either one of 2 states: X_t^+ (having an associated reward z_t^+) with probability p_t^+ , or X_t^- (having an associated reward z_t^-), where $p_t^+ + p_t^- = 1$. In the sequential auction case considered here, X_t^+ , respectively X_t^- represent the states in which the agent wins / does not win an upcoming auction (the formal link is made in Sect. 2). The following relation holds:

$$CE[(CE[(z_{t+1}^+, p_{t+1}^+), (z_{t+1}^-, p_{t+1}^-)], p_t^+), (z_t^-, p_t^-)] = CE[(z_{t+1}^+, p_t^+ p_{t+1}^+), (z_{t+1}^-, p_t^+ p_{t+1}^-), (z_t^-, p_t^-)]$$

Note that the notations z^+ and z^- , for each time t only relate to whether the agent wins or does not win the auction. These numbers can actually be negative, in case the payoff for a state is negative. For example, an agent with a strictly complementary valuation for items sold at times t and $t+1$, that wins the item at time t but does not win the item at time $t+1$, gets the payoff z_{t+1}^- , which is negative.

Proof: The proof involves repeated application of Eq. (4.4) to the left-side term:

$$\begin{aligned} CE[(CE[(z_{t+1}^+, p_{t+1}^+), (z_{t+1}^-, p_{t+1}^-)], p_t^+), (z_t^-, p_t^-)] &= \\ &= -\frac{1}{r} \ln[p_t^+ e^{-rCE[(z_{t+1}^+, p_{t+1}^+), (z_{t+1}^-, p_{t+1}^-)]} + p_t^- e^{-rz_t^-}] \\ &= -\frac{1}{r} \ln[p_t^+ e^{-r[-\frac{1}{r} \ln[p_{t+1}^+ e^{-rz_{t+1}^+} + p_{t+1}^- e^{-rz_{t+1}^-}]]} + p_t^- e^{-rz_t^-}] \end{aligned}$$

After reducing $-r(-\frac{1}{r})$ and using that $e^{\ln X} = X$, we get:

$$\begin{aligned} &= -\frac{1}{r} \ln[p_t^+ p_{t+1}^+ e^{-rz_{t+1}^+} + p_t^+ p_{t+1}^- e^{-rz_{t+1}^-} + p_t^- e^{-rz_t^-}] \\ &= CE[(z_{t+1}^+, p_t^+ p_{t+1}^+), (z_{t+1}^-, p_t^+ p_{t+1}^-), (z_t^-, p_t^-)] \text{ (q.e.d.)} \end{aligned}$$

Note that the above property can be applied recursively to games with any number of stages. This property, while apparently straightforward, is important since it shows that performing local CE optimization at each time step gives the same result as CE optimization for the entire game (a property which is not obvious for non-linear functions). As such, it is used as an implicit assumption in our MDP model.

4.2.1 The importance of risk aversion in decision making: an example

In the following, we give an illustration why risk aversion can have an important effect on monetary values. Consider the case of two complementary-valued items: A and B, which are sold sequentially. Suppose the agent has to accept a sunk cost of \$5 (dollars or any monetary units) for item A. If she acquires both A and B, she makes a profit of \$10, but if she doesn't, she makes a loss of -\$5 (thus potential profit is double the size of potential loss). Supposing the agent estimates the probability of acquiring B at p_B , how large does p_B have to be in order for the agent to accept the gamble?




Figure 4.1: Example of the certainty equivalents of 3 agents with 3 different risk profiles for a lottery with 2 possible outcomes: -\$5 (non-desirable) and \$10 (desirable). The figure illustrates 2 cases: A(left): The desirable outcome is assigned a monetary value of \$10 by all agents. B(right): The desirable outcome is assigned a monetary value of \$5 (for the risk indifferent agent ($r \rightarrow 0$), \$7.5 by the slightly risk averse one ($r=0.15$) and \$10 by the strongly risk-averse agent ($r=0.3$))

We plot the CE payoffs in this lottery for 3 risk attitudes of the agents, from $r \rightarrow 0$, $r = 0.15$ and $r = 0.3$. The left-hand side of Fig. 4.1 shows the case when all agents have the same evaluation for both the desirable (i.e. +\$10) and the non-desirable (-\$5) outcome. From this figure, one can already see that a risk neutral agent ($r = 0$) would “join in” this

lottery or sequence of auctions, if the probability of winning (getting the desirable outcome) exceeds 33.3%. However, a relatively risk-averse agent ($r = 0.3$) would need to have at least 78% probability of winning in order for it to assign a positive CE value to this lottery (and thus have an incentive to participate in the game). In the right-hand side of Fig. 4.1, we keep the payoff of the non-desirable outcome constant at $-\$5$, but we vary the maximal payoff from $\$5$ (for the risk indifferent agent), to $\$7.5$ (for $r = 0.15$) and $\$10$ (for $r = 0.3$). Even if the estimated probability of acquiring the bundle $\{A, B\}$ is exactly the same for all 3 agents, the probability of winning has to be above 97% in order for the agent with the highest valuation to assign the sequence of auctions the highest CE value, among these agents.

4.3 Bidding in sequential auctions with complementarities

As shown in the introduction, the main problem that a bidder has to face in a sequential auction with complementarities is the exposure problem. Following Boutilier et. al. [27] and Greenwald & Boyan [89], we model the decision problem that the bidder agent has to face in sequential auctions as a Markov Decision Process.

Assume there is a set of items I_t , sold in sequential auctions held at time points $t = 1..n$. A state in this game is specified by a set of goods X_t acquired up to time t (where $X_t \subseteq I$ for $t = 1..n$). The *bidding policy* of an agent in this game is described by a vector of bids $\vec{B} = (b_1, \dots, b_n)$, which assigns a bid b_t to each item sold at time point t . Fig. 4.2 illustrates this, for an auction with 2 items.

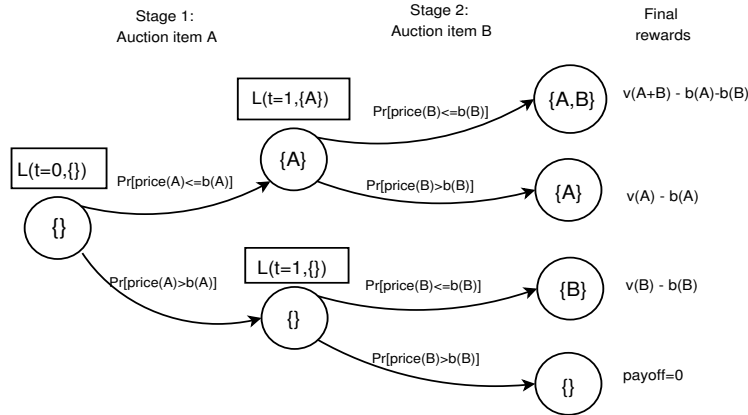


Figure 4.2: The decision process faced by an agent in sequential auction, for a two stage example, with goods labeled A and B

The bidding agent maintains a probabilistic expectation of the closing prices for items I_1, \dots, I_n , in the form of n distributions. In the current model, these distributions are assumed independent of each other and stationary during one bidding round of n auctions (n could

also be seen as the number of auctions the agent can stay in the game before its deadline). This definition of stationarity does not exclude the agent being able to learn, or refine its distributions of closing prices between episodes but, in this chapter, we assume they are stationary for the duration of n auctions (i.e. one episode).

Considering the probabilistic distribution of future prices (a similar choice as in [1, 27, 89]) is more relevant to this setting than simply working with a vector of the average past prices (such as in [184, 217]), since the thickness of the tails of the distribution may be of particular importance if the agents are risk averse. Note that in this form, we do not make any assumption on the type or shape of the expected future distributions: they can be normal, log-normal (usually used to model future prices in financial markets), uniform, binomial etc. For the results reported in this chapter, we employed the normal distribution, but the generic approach can be applied to other distributions as well. The transition probabilities between different states are the cumulative distribution probabilities that the agent wins the lottery with its current bid b_t :

$$Prob(X_{t+1} = X_t \cup \{I_t\}) = Prob(ClosingPrice_t \leq b_t) = cdf_t(b_t)$$

where $cdf_t(b_t)$ denotes the *cumulative density function* of the probability distribution over the closing prices, when bid b_t is placed.

We model the utility of a future outcome at each time step t (except the final one when all the goods have been allocated) as equivalent to a lottery $L_t(X_t, b_t)$. The payoffs of this lottery are determined by the agent's utility function, the set of items acquired so far X_t and bid b_t . The probabilities over outcomes depend on the bid b_t and expectation of future price distributions. The decision problem the agent faces, at each time point is to choose a bid b_t that provides the right balance between expected payoff and probability of winning, given her risk aversion r . This means choosing b_t which maximizes the certainty equivalent of lottery $CE(L_t(X_t, b_t))$. Using formal MDP notation, the value at each state is:

$$Q(X_t, b_t) = CE(L_t(X_t, b_t))$$

The optimal bidding policy and the corresponding reward as:

$$b_t^* = \pi(X_t) = \operatorname{argmax}_{b_t} Q(X_t, b_t)$$

$$V(X_t) = \max_{b_t} Q(X_t, b_t)$$

We can rewrite the above two equations, for the optimal bid at time t b_t^* and the associated optimal certainty equivalent value CE^* , that can be obtained by taking the optimal bidding decision as:

$$b_t^* = \operatorname{argmax}_{b_t} CE(L_t(X_t, b_t))$$

$$CE^*(L_t) = \max_{b_t} CE(L_t(X_t, b_t))$$

Note that this optimization of the certainty equivalent value CE is performed for the current auction at time t , but assuming that the optimal bidding decisions are taken the

whole sequence of future auctions, occurring at times $t + 1, \dots, n$. Therefore, the problem of determining the optimal bid b_t^* for time point t actually involves recursively determining the bids b_t^*, \dots, b_n^* that maximize the certainty equivalents of states at t_1, \dots, n . Due to the recursive property of the CE function (captured by Lemma 1 above), maximizing $CE(L_t)$ at each state leads to maximizing the initial certainty equivalent expectation for the entire sequence of auctions, i.e. maximizing $CE(L_0)$. This means that standing MDP reasoning models can be applied to this problem, where the Q values of the standard MDP definition are the CE values of the lottery over future expectations at each step.

A naive alternative to this method would be the application MDP optimization directly to the utility function of the agent (as done in [27] for risk neutral agents). For risk-averse agents, however, due to the non-linear nature of the utility functions, definitions of bidding policies in sequential auctions can only be defined in terms of the CE values of future states³. This is done in the following Sections, which also include a numerical example and an illustration that provides insight into the dynamics of the problem.

4.3.1 Optimal bidding policy for sequential 2nd price (Vickrey) auctions

Greenwald & Boyan [89] show that the optimal bidding strategy for a risk-neutral agent in a second-price sequential auction is to bid the difference between the expected value of the state when the auction is won and the expected value of the state when the auction is not won. Here we can extend these results to the risk-averse case as follows.

Suppose at time t (after a set of t previous auctions) the agent is in a state in which he has the set of items X_t . At the next step (i.e. after the auction occurring at t), he can transition in either one of two possible states: one in which he obtains the set of items $X_{t+1}^+ = X_t \cup \{I_{t+1}\}$ (if the auction is won) or $X_{t+1}^- = X_t$ (if the auction is not won). If the auction at time t is a second-price one, the optimal bidding policy available to the agent is:

$$b_t^* = CE(L_{t+1}(X_{t+1}^+)) - CE(L_{t+1}(X_{t+1}^-)) \quad (4.5)$$

assuming that at all subsequent steps $t + 1, \dots, n$ the locally optimal bids are chosen.

Proof 3 *The proof resembles the proof in the textbook of Krishna [130], which refers, however, only to risk-neutral bidders. First, we simplify the notation by denoting the certainty equivalent of the state when the item is acquired by $CE_{t+1}^+ = CE(L_{t+1}(X_{t+1}^+))$ and the certainty equivalent of the state when the item is not acquired by $CE_{t+1}^- = CE(L_{t+1}(X_{t+1}^-))$. There is a set of n_a independent bidders in each auction, that only desire the item sold in that auction. All auctions being second price, they always have a dominant policy of*

³We stress that the term “optimal” used in this chapter, should be interpreted as optimal w.r.t. the bidder’s aversion to risk and estimation of future price distributions. This is not the same concept as dominant bidding strategy from standard auction theory (i.e. independent of the behaviour of other bidders). As discussed in the introduction, dominant strategies are not known to exist for the sequential settings considered in this paper.

bidding their true value. Let the valuation function over these n_a bidders be denoted by $G(x) = G_i(x)^{n_a}$, where $G(x)$ is the cumulative distribution that bids of all n_a agents are smaller than x ($G_i(x)$ here refers to a single independent bidder, but we can consider them in aggregate, without loss of precision). Then $g(x)$ is the density function of this distribution, i.e. it denotes the probability that the highest bid of the independent bidders is exactly x .

Note that the state of winning the auction and having to pay x brings a monetary gain of $CE_{t+1}^+ - x$ for the agent, while loosing brings a monetary gain of CE_{t+1}^- . In this case, however, the amount to be paid depends on the highest bid of independent bidders, so the standard certainty equivalence definition needs adjusting. Basically, the CE of bidding b_t in a state at time t can be expressed as:

$$CE(b_t) = -\frac{1}{r} \ln \left\{ \int_0^{b_t} g(x) e^{-r(CE_{t+1}^+ - x)} dx + (1 - G(b_t)) e^{-rCE_{t+1}^-} \right\}$$

The optimal bid b_t^* can be obtained by taking the derivative of the above expression, i.e. when $\frac{dCE(b_t)}{db_t} = 0$. This gives:

$$\frac{-g(b_t)}{r \int_0^{b_t} g(x) e^{-r(CE_{t+1}^+ - x)} dx + (1 - G(b_t)) e^{-rCE_{t+1}^-}} (e^{-r(CE_{t+1}^+ - b_t)} - e^{-rCE_{t+1}^-}) = 0$$

Of this expression, the first fraction is never zero and can be reduced, which basically gives:

$$e^{-r(CE_{t+1}^+ - b_t)} = e^{-rCE_{t+1}^-}$$

Which finally, after applying the logarithm and dividing by $(-r)$ gives:

$$CE_{t+1}^+ - b_t = CE_{t+1}^-$$

Resulting in the final expression for b_t as:

$$b_t^* = CE_{t+1}^+ - CE_{t+1}^-$$

So basically, the marginal optimal bid b_t^* in a sequence of second price auctions is always the marginal difference between the certainty equivalents of the next two states. Note that this was known from standard auction theory for the case of risk-neutral bidders [130]. Basically, since a rational agent views all previous payments as sunk costs, they can be discounted and do not have to be accounted for in future bids. The intuitive reason why we find this result in the case of risk-averse bidders as well is that certainty equivalent functions, although not linear, are basically monotonically increasing in the monetary payoffs of future states, so it is rational for the agent to increase her bid until the difference $CE_{t+1}^+ - CE_{t+1}^-$ is covered.

4.3.2 Optimal bidding policy for sequential 1st price auctions: numerical solutions

For first-price auctions no closed form optimal bidding policy can be formulated, because agents have, as in the case of risk-neutral agents, an incentive to shade their bid. Liu et. al. '03 [146] show, for the case of single-shot first-price auctions that, on average, risk averse agents shade their bid less than risk neutral agents, since they want to minimize the chance of losing the auction. In this case, the optimal bid level b_t^* given in Eq. 4.5 above for the second price auction represents an upper bound on the bid level a rational agent would place in a first-price auction.

For the sequential case, in order to get insight into the case, we computed the numerical solutions of the optimal bidding policy as perceived by the agents at time $t = 0$ (before entering the sequence of auctions). This is done for a sequence of 2, respectively 3 upcoming auctions (items are numbered alphabetically, by the order they are being auctioned). The analysis can be extended to any number of auctions, and the results are largely similar).

We take the expected distributions for future prices for individual items are drawn from identical, independent *normal* distributions (i.i.d.s are a choice widely used in economic modeling [153]). In this case, we chose normal distributions with mean $\mu = 2.5$ and dispersion $\sigma = 1.5$. The chosen valuations levels are: $v_{\{A\}} = 0$, $v_{\{B\}} = 0$ and $v_{\{A,B\}} = 10$ (for the 2-stage auction), respectively $v_{\{A,B,C\}} = 15$ and 0 for all other subsets (for the 3-stage auction). This choice of values is such that the sum of the mean expectation of the costs is exactly half the bundle payoff.

A bidding policy is defined as a combination of bids for item b_A, b_B , with the note that the bid for B is only placed if the agent wins A in the preceding auction (otherwise, it has a dominant policy to bid 0 and earns a reward of 0). Using a mathematical optimization package (in our case Matlab), we computed the optimal bid levels of this game (b_a^*, b_B^*), for each level of risk aversion from 0 to 1, as well as the expected CE level of this optimal bidding policy, i.e. $\max_{b_A, b_B} CE(L_{t=0})$.

In Fig. 4.3, we show the CE value of the initial choice to enter the set of auctions (i.e. $CE(L_{t=0})$) for one level of risk aversion r and all possible combinations of bids for the first, respectively second good in the sequence. As can be seen in Fig. 4.3, the surface of possible bids has a single optimum point, for each level of risk aversion.

In Fig. 4.4 we plot the optimal bid levels for a sequence of 2, respectively 3 auctions. Basically, each point on the left (i.e. two-item) side of Fig. 4.4 corresponds to the coordinates of the optimum point in exactly one bidding surface, such as shown in Fig. 4.3. The same can be said about the right side (i.e. the 3 item case), although in this case the bidding surface cannot be actually visualized (being 4-dimensional).

From the analysis of Fig. 4.4, we can already highlight some important effects:

- The more averse a risk agent is, the higher she will bid for the second item in a 2-stage auction sequence. Intuitively, a risk averse agent is more concerned with reducing as

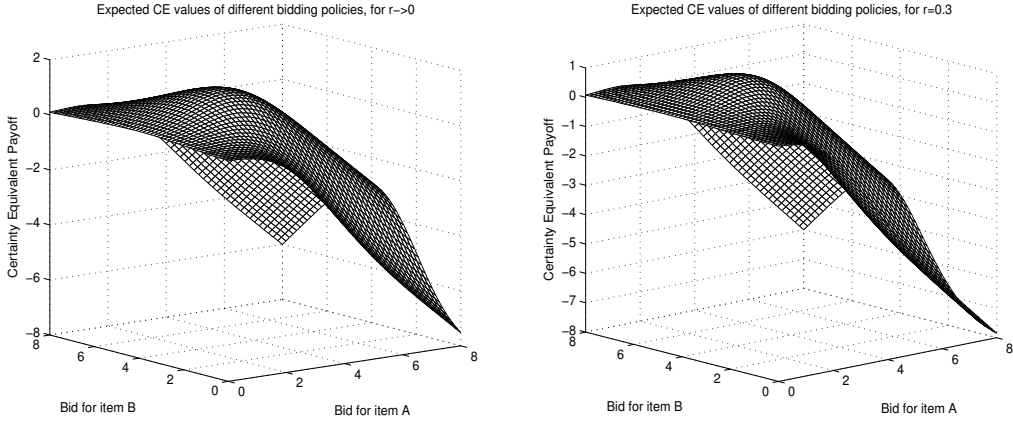


Figure 4.3: Example of the certainty equivalent payoff in a two-stage sequential auction for 2 items: A (at time $t=1$) and B (at $t=2$). The graph shows the CE value of the corresponding 2-stage game, if the costs for both items are drawn from $N(\mu = 2.5, \sigma = 1.5)$, for an agent with $r \rightarrow 0$ (left) and $r = 0.3$ (right).

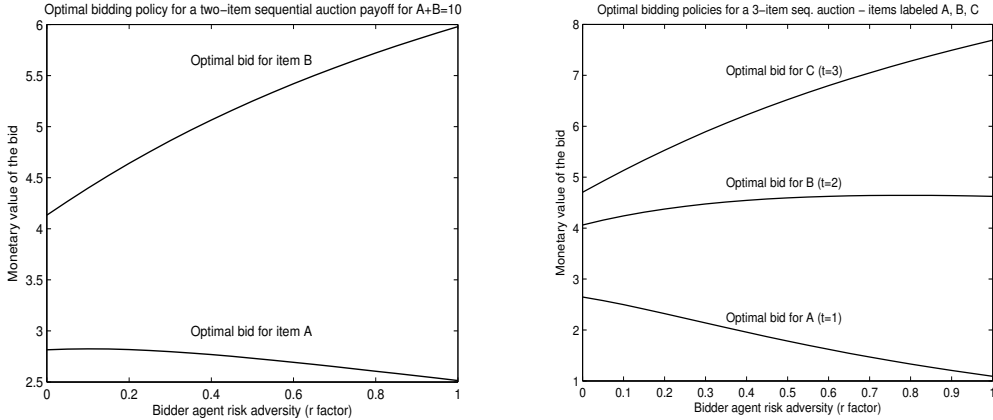


Figure 4.4: The optimal bidding policy available to an agent having risk aversion r , in a 2, respectively 3-stage sequential auction. The items have a complementarity value of \$10 (resp. \$15) if acquired together, but no value if acquired separately. The costs for all items are drawn from a normal distribution $N(\mu = 2.5, \sigma = 1.5)$.

much as possible the probability she will loose the auction for B and not cover her sunk cost for item A. By contrast, a more risk-neutral agent is willing to accept a slightly higher probability she will have a sunk cost, if the potential gain is greater. Otherwise stated, agents with different risk profiles have different levels of awareness

of costs already incurred.

- By contrast, the optimal bid level for item A slightly decreases as the agent becomes more risk averse. Risk averse agents are not willing to accept a high sunk cost - thus their optimal policy is to avoid bidding aggressively in the first round. They may prefer not to participate at all in the sequence of auctions, than to win the first auction with a high sunk cost, which would be difficult to cover. Furthermore, note that in this example, the average mean expectation of cost of the first item is only a quarter (\$2.5) of the maximal possible cost. We also performed tests with other mean expectation costs, and found that, if these costs become higher, the effect is considerably more pronounced - and risk-averse agents' optimal bid policy may simply be not to participate at all in the auction sequence.

4.3.3 Bidding strategy for multiple copy auction sequences

The MDP-based bidding strategy outlined above can lead to an optimal bidding policy, but only if all CE values of the states for the entire game are computed. This can become computationally expensive, especially if the sequence contains many stages (auctions). In the simulations presented in Sect. 4.4 below, we make an approximation that enable us to significantly prune the state tree in solving the multiple copies problem. This problem appears when the bidding agent is interested in only a limited number of items to form a useful bundle, but these are offered for sale repeatedly.⁴ Suppose items are divided into several types. The agent's expectation of closing price distributions for all items of a given particular type is the same (thus she does not model the future expectation probability per auction or per item, but per type of item). If this expectation remains the same during the number of bidding rounds the agent stays in the game, then it is possible to reduce the state tree representation from a representation dependent on the number of future auctions to a representation which depends only on the size of the bundle the agent wishes to buy.

Formally, if there are several items of type A and the agent knows that there are n_A more auctions of items of type A to take place. Then the probability of transition from any state X to a state $X \cup \{A\}$ (i.e. winning at least one item of type A at some point in the next sequence of n_A opportunities), given that the agents bids b_A in each of the auctions in that sequence is:

$$Prob(ClosingPrice_A \leq b_A) = 1 - [1 - cdf_A(b_A)]^{n_A}$$

The above formula can be used to determine the probabilities of the getting an item of type A in the final state (i.e. after all auctions for a good of type A have closed). One still needs to apply the MDP to determine the best policy based on these probabilities, but this is straightforward, as it does not require computing the whole tree.

Note that this policy only uses as input the number of future auctions of each type remaining before the auction has to leave the market, not their exact order. In fact, if one

⁴Multiple copies can be seen as an instance of the substitutability problem - though substitutability is wider, if we allow for partial substitutes. These are not considered in the current work.

knows the exact order that future auctions take place in, then it might be better to compute the whole tree (although that's exponentially more expensive). However, not knowing the exact order that future auctions will take place in, only the number of auctions of each type, is more realistic in many real-world settings. So, for example, in the simplified transportation case shown in Section 4.4.5 below, in practice, planners may know that a number of opportunities (i.e. transportation orders) to fill a truck may appear before the truck needs to start driving, but they don't know exactly the order in which these will be offered.

For the experiments reported in this chapter, because goods are all of the same type (even if a bidder may desire only a bundle consisting of several such goods), this heuristic approximates very well the optimal bidding policy. In this case (i.e. same-type goods), there is basically only one possible sequence of future auctions, and the length of this sequence basically represents the full information needed to describe it. If there are several possible types of bundles, then the difference in performance may depend on the exact auction sequence. However, even this problem can be mitigated by randomizing over all possible auction sequences when performing experimental evaluation.

As we discussed in the numerical example, having multiple future opportunities to buy a good may determine risk-neutral agents to reduce their bids (since there is a higher chance of winning one of them), but it may also encourage risk-averse bidders to join the bidding, bidders which would otherwise find a short sequence of auctions to be too risky to participate.

4.4 Experimental analysis

The goal of the experimental results presented in this chapter is to test how different sequential auction market settings are influenced by the presence of a complementary valuation bidders, with different risk aversion levels. We look at how risk aversion influences the expected profit that the synergy bidder makes over a sequence of auctions, as well as the probability of completing the desired bundle and ending up with an incomplete bundle (which can result in a loss). Furthermore, we also study how the expected revenue of the seller(s) is influenced by the presence of a synergy bidder in a market, as well as how the number of buying opportunities (i.e. length of the auction sequence) influences the expected profits of seller and buyer.

The first part of this chapter studies these questions for a market with a single synergy bidder participates in a sequence of auctions for items of the same type. More concretely, we assume a market consisting of a sequence of auctions, each populated by a set of local (single-item) bidders and one synergy (global) bidder that desires exactly one bundle of two items. The number of auctions that the synergy bidder can stay in the market is fixed for each simulation round (although this parameter will be varied between different experiments). In the second part of our experimental study, we introduce bundle differentiation, i.e. the auction sequence consisting of auctions for two types of items and a synergy buyer that can choose between the two possible bundles. This setting was motivated by a transportation logistics setting described in Section 4.4.5.

4.4.1 Experimental hypotheses

In order to better structure the presentation, we first formulate three hypotheses, that should be confirmed or disproved in the experimental tests. These hypotheses are intuitively formulated based on the properties observed in the theoretical part of this chapter and should help the reader understand better the focus and choices made in the simulation model.

Hypothesis 1: A more risk averse agent will have a lower chance of ending the sequence of auctions with an incomplete bundle (i.e. a bundle in which the first item is acquired, but not the subsequent ones, hence resulting in a loss).

Note that the statement in Hypothesis 2 appears obvious: it is more a control hypothesis. If we do not find this, then there may be reasons to believe something is wrong in our experimental set-up. The most important side effect is stated as:

Hypothesis 2: A synergy bidder with a higher risk aversion will obtain a lower average profit from bidding in a sequence of auctions than a synergy bidder which is less risk averse.

The final hypothesis refers to the case of different auction lengths.

Hypothesis 3: For all risk aversion levels, the expected profit of a synergy buyer desiring a bundle of items will be higher if there are more auctions in the sequence (i.e. more opportunities to buy), while the chance of ending up with an incomplete bundle will be lower.

Besides these hypotheses, referring to the buyer, in our experiments we also look at the average revenue that a seller of a set of items sold in sequence will be lower if the synergy buyer present in the market is more risk averse. There are a further two further hypotheses, related to markets with different item types, but they will be introduced them later.

4.4.2 Experimental setup

The experimental set-up used is as follows. We consider a sequence of n closed, first-price auctions, in all of which exactly one item of the same type A is sold. In each of these auctions, there are an (unspecified) number of local bidders, assumed myopic, that desire exactly one item of type A . Since these agents are assumed myopic (i.e. they only consider the current auction they participate in), we can model their bids in each auction through some random distribution. Note that the myopicity assumption of local bidders is important here: if the bidders are able to strategize over the sequence of auctions, or over the presence of a synergy buyer in this sequence, then the model we use for their bidding behaviour may not hold. In this model, because in all auctions an identical good is sold, we can model

the maximum bid received from the competition in each of the n auctions in the sequence through identical, independent probability distributions (i.i.d.) - a choice that is also made in other decision-theoretic bidding models, e.g. [83]. Since we do assume any prior information about the way independent bidders place their bids, we take the most general case and assume they follow normal distributions $N(\mu, \sigma)$.

In each sequence of auctions there is exactly one synergy (or global) bidder participating. This bidder desires exactly one bundle of two items of the same type A (and is assumed to have no value for an individual item). The synergy bidder must acquire this bundle in exactly n auctions (here, the number n of auctions that a bidder can stay in the game can also be thought of as a way to model a shorter or longer deadline that a bidder has). The value $v(2 * A)$ that the synergy bidder assigns to the useful bundle, the number of auctions n , as well as the parameters μ and σ that model the behaviour of single-item bidders are all parameters of the simulation. For each market configuration, average results are reported over 1000 runs.

4.4.3 Experimental results for one-type item auctions

In this section we give the results for a market setting where only one type of item A is sold, and the synergy bidder desires exactly one bundle of two such items. Initially, we do this for sequences of $n = 7$ auctions, where the values of independent bidders are drawn from a distribution $N(\mu = 4, \sigma = 2)$. The synergy bidder assigns a value of $v(A, A) = 10$ (notice that this means the synergy buyer values a bundle of two item, on average, with 25% more than the independent bidders). The bids of the synergy buyer are computed according to the heuristic in Sect. 4.3.3, based on the risk aversion coefficient shown on the abscissa. Results (with averages over 1000 runs) are shown in Fig. 4.5.

Returning to the hypotheses stated in Sect. 4.4.1, we see that indeed, Hypothesis 2 is confirmed for this setting: the higher the risk aversion of a synergy bidder, the lower his/her average expected profit (the drop is quite considerable - from 4 to around 2.5). However, one can also notice the variance of the results decreases slightly for the risk-averse bidder. These results are consistent with expectations.

The right side of Fig. 4.5 shows the average revenues of the seller for this setting. These are somewhat surprising, given that one would intuitively expect seller revenues to drop if bidders in the market are more risk averse. In fact, it seems there is some type of revenue equivalence, in the sense that the revenue that a seller can expect in such a sequence of sequential auctions with a risk averse bidder does not depend on his/her risk aversion. From further examination, this can be explained as follows: since these are first price auctions, a more risk averse bidder will bid less often in this sequence. However when he does bid, he will bid considerably more which, on average, has a compensating effect for the reduced participation. Nevertheless, a more in-depth investigation is needed (with a larger market and more synergy bidders), in order to formulate a hypothesis regarding this point.

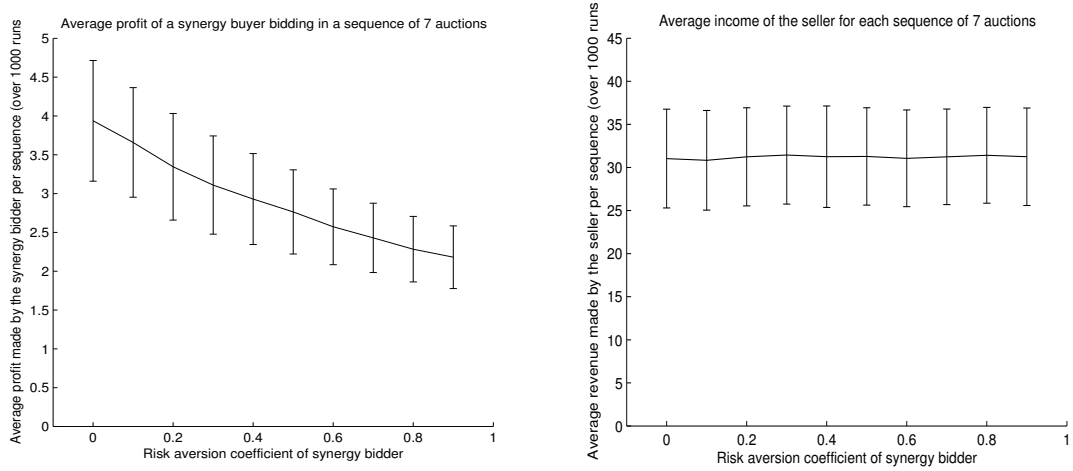


Figure 4.5: Setting with 7 sequential auctions and one synergy buyer with $v(A, A) = 10$, and independent agents for each auction, bidding according to $N(4, 2)$. The left side give the average profit of the synergy buyer, while the right the average revenue of the seller (both averaged over 1000 runs)

4.4.4 Results for one item and different auction lengths

In this Section, we extend the above analysis to a setting where we vary not only the risk aversion r of the synergy buyer, but also the number of auctions he/she can participate in to acquire the desired two-item bundle. Furthermore, we made this setting more competitive: while the valuation of each synergy buyer for a bundle of 2 A-s remains $v(A, A) = 10$, the competition is slightly more aggressive and will be according to $N(4.5, 2)$. This basically means that we reduce from 20% to 10% the advantage in valuation that a synergy buyer has, on average, over the independent bidders. The motivation for this is that it seems more relevant to study how the success rate is influenced by multiple buying opportunities (i.e. varying number of auctions), for a more competitive setting.




Figure 4.6: Results for the average profit of the synergy buyer (left) and seller revenue (right), for a setting with one synergy agent with $v(2A) = 10$ and a set of independent agents bidding according to $N(4.5, 2)$. The number of auctions the synergy agent can stay in the game to acquired the desired bundle, as well as his/her risk aversion coefficient are varied for the different settings. All results are averages over 1000 runs, but to avoid overloading the picture, error bars were not included.




Figure 4.7: Success and failure rates for a setting with one synergy agent with $v(2A) = 10$ and a set of independent agents bidding according to $N(4.5, 2)$. The left side graph show the percentage (among the 1000 runs) in which the agent acquired his/her target bundle of two A items. The right side graph shows the percentage of runs with incomplete bundles (i.e. runs in which the synergy buyer obtained the first item, but not the second, hence resulting in a loss).

Results from these tests are shown in Fig. 4.6, respectively 4.7. Returning to the above stated hypotheses, we can conclude that, indeed Hypothesis 2 can be confirmed. Even for this modified setting, there is a marked decrease in average synergy bidder profits, as his/her risk aversion increases. This effect is more noticeable for the higher auction lengths (5, 7 or 10 items). The reason for this is that, for this competitive setting, the participation rates, even for the more risk-neutral agents are on the low side (see left side of Fig. 4.7) and relatively constant over r . Furthermore, it seems that the revenue the seller can expect is also relatively constant over the risk aversion of the synergy buyer, even for these more general tests.

From looking at Figs. 4.6 and 4.7 one can clearly see the very large effect that the number of available auctions has, both on the expected profit and success rate of the synergy buyer (thus confirming Hypothesis 3 above). This effect clearly holds for all risk aversion coefficients of the synergy buyer and all configurations. From Fig. 4.7, one can also see that a more risk averse bidder, while making, on average, less profit, does have some advantages in this type of auctions: his/her chances of ending up with an incomplete bundle before the deadline (hence making a loss in that particular auction run) decrease considerably. Thus, these results support Hypothesis 1, as expected. Since in many real life bidding situations (one will be discussed in the following Section), agent consider only the possibility of profit/loss in an immediate run (not the long-term statistical average), minimizing chance of a loss, even if it has only 5%-10% probability, can be an important consideration.

4.4.5 Setting with different item types and more complex preferences

The previous Section has already highlighted the complexity of bidding in sequential auction to get a bundle of two items, even for the simplified setting with one possible item type. However, in most real-life scenarios, on top of the question of how to divide their bids between complementary items in a sequence, agents are confronted with several alternatives that they must choose from during bidding. In fact, the potential complexity of the space possible preferences is very large. In this Section, while we do not completely model the full potential complexity of possible preferences, we show that having a second type of good to choose from introduces a whole different dimension to the dynamics of decision-theoretic bidding in sequential auctions.

We should mention that our choice for the valuation structure of the bundles, while simple, is motivated by a transportation logistics application setting and does capture much of the dynamics of that use case. Therefore, before we describe the experimental set-up and results, we motivate it by briefly describing how the experimental choices made could plausibly fit a real-life application setting.

Bidding in sequential auctions for transportation orders

The problem setting we considered in our auction model is that of distributed transportation logistics with partial truck loads (a real-life, business-oriented platform for this case,

developed in collaboration with a large logistic company, is described in Chapter A of this thesis).

In the logistic setting we consider, transportation orders (either from one, but usually from different sellers/shippers) are usually sold at different points in time through spot market type mechanisms (usually auctions). The bidders for these loads are small transportation companies who try to acquire a suitable set (bundle) of orders that would fit the capacity of their trucks. In this model, we assume all orders are ready for pick-up or return delivery at one central transportation depot⁵. Fig. 4.8 shows just such a topology, with delivery point group into 2 main delivery regions).




Figure 4.8: Example transportation scenario with one central depot D and two disjoint transportation regions: A and B.

Acquiring suitable combinations (bundles) of orders to fit the same trip with one truck is crucial for profitability in this setting. A truck acquiring, for example, an order for 1/2 truckload to be delivered to a certain region usually counts on acquiring another 1/2 truckload order from the same region, in order to make a profit. In this case, item types represent different delivery regions - each trucking company expecting different costs/profit structure per region, depending on its transportation network. Another possibility for bundling can concern symmetrical outgoing/return orders which originate in the same region.

In the utility model used in our auction simulations, we abstract the main characteristics of this setting. In this way, bidders can be considered as truck owners (i.e. carriers), the items are transportation orders, item types correspond to different delivery or pick-up regions. In practice, auctions for transportation orders are reverse auctions: the bidders that offer the lowest cost get the order. However, the corresponding model with sequential ascending auctions studied in this paper is basically equivalent to this, and it's easier to compare with other models and in existing literature.

Furthermore, in reporting the results, we also make the assumption that there is a single seller for all the goods (or orders) in the sequence. While in practice transportation orders may originate from multiple shippers (or customers), the aggregate revenue of the single seller can be seen as indicative of a global average, that a seller, without knowing his/her precise place in the sequence of auctions, has from selling items in this sequence.

⁵This is actually a realistic assumption in many cases, especially if there is just one shipper, or several small shippers who aggregate their demand to one central distribution point.

Multiple item simulation set-up

The sequential model we consider is as follows. The number of auction rounds is still fixed at 7, but there are two types of goods: A and B. In this setting, we introduce a differentiation between the items: items of type B are relatively “rarer” (they are sold only in 2 auctions out of the 7), while items of type A are more common, and sold in 5 out of 7 auctions.

However, the value the synergy buyer assigns to a bundle of such items is also asymmetrical. A bundle of 2 items of type B has a valuation of $v(B, B) = 20$, while a bundle of 2 items of type A: $v(A, A) = 10$. The competition coming from single-item bidders for those goods is also different. For goods of type A, the bids from this competition are modeled through a normal distribution $N(\mu_A = 4, \sigma_A = 2)$, while for items of type B through a distribution: $N(\mu_B = 6, \sigma_B = 2)$. Therefore, the additional valuation of the synergy agent for a bundle of two items, compared to the *average* paid by independent bidders is only $2/10 = 20\%$ for a bundle of type A, but $8/20 = 40\%$ for a bundle of type B. at the same time, a bundle of type B is twice as rare.

4.4.6 Multiple item setting: hypotheses

Before we present the result graphs for this setting, we follow the format of the previous Sections, and formulate two additional hypotheses:

Hypothesis 4: In a market with two types of items, one of which is rarer, but also more valuable than the other, the synergy bidders with a risk aversion coefficient above a certain level may select to bid for a bundle of the more common item, in order to maximize their chances of completing the bundle. This can reduce the synergy bidder’s average expected profit from the auction sequence.

Hypothesis 5: In the above setting, if risk averse agents prefer to bid for more common, but less valuable goods, this also reduces the revenues of the auctioneer.




Figure 4.9: Results for the average profit of the synergy buyer (left) and seller revenue (right), for a setting with two items A (of value $v_{A,A} = 10$) sold in 5 auctions, and B (of value $v_{B,B} = 20$) sold in 2 auctions. Notice there is a transition, because agents with risk aversion $r \geq 0.5$ do not try to get the higher value bundle (of item B).

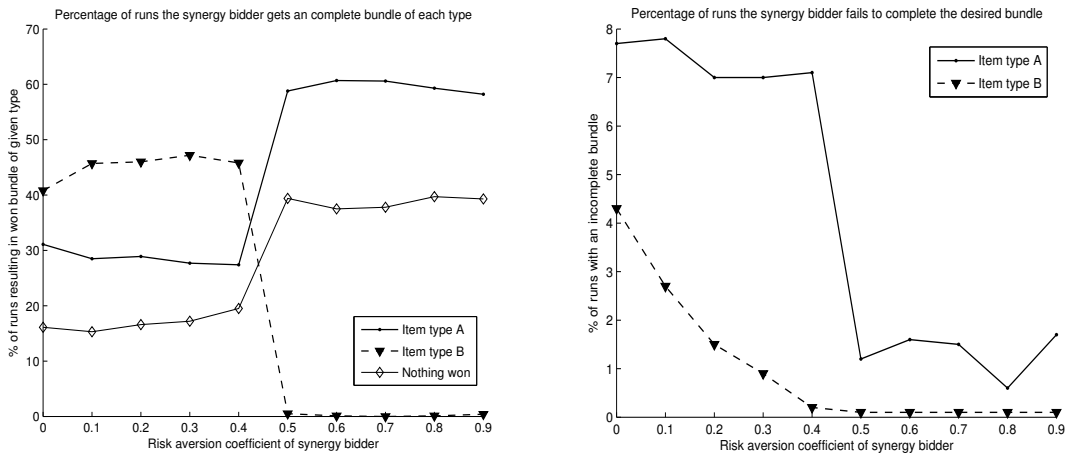


Figure 4.10: Percentages of success and failure per 1000 simulation runs, for a setting with two different type of items described above. Notice the transition at $r \geq 0.5$, showing that risk averse bidders do not try to get the bundle with the rarer item B, but only one of A.

4.4.7 Results for two-item case

Experimental results for the above setting are presented in Figures 4.9 and 4.10. Fig. 4.9 gives the average profit of the synergy agent and the seller, while Fig. 4.10 gives the percentages of auctions that agents complete (or fail to complete) bundles of items of either type. All results reported are averages over 1000 runs.

Basically, Hypotheses 4 and 5 are, on the whole, confirmed by these tests: there is a decrease in the expected profit of the synergy buyer and seller. However, there is an important caveat: there seems, for this parameter settings, to be an important threshold effect as the risk aversion factor of the agent becomes $r = 0.5$. The reason for this threshold effect is clear from Fig. 4.10: at this level, the more risk averse agents stop trying to bid for the more valuable, but also “riskier” bundle of item B (for which there are only 2 available auctions), and go for a bundle of item A, from which there is less absolute profit to be made, but for which there are 5 available auctions.

The left-hand side of Fig. 4.9 shows that, while going for the bundle of item B brings, on average, slightly more profit, this result also is subject to a much higher variance, i.e. the bidding agents are more likely to loose money by failing to complete their desired bundle. By contrast, bidding for a bundle of type A (as the more risk-averse agents do), can slightly decrease the average expected profit, but the bidder is less likely to loose money. In fact, the lower interval of the variance bars, in this case, seem to be all above the zero axis.

The seller revenue (right side of Fig. 4.9) is also influenced by the risk aversion (and, hence, the bidding behaviour) of the synergy bidder, but the decrease in seller revenue that occurs at the threshold level is relatively slight (of only a few percentage points). Nevertheless, one should note this is a setting with only one synergy bidder present in the market, therefore the average effect may be understated.

4.5 Conclusions and further work

To summarize, the main contributions of the work presented in this Chapter are as follows.

First, we establish a formal link between bidding strategies in sequential auctions and standard (Arrow-Pratt) risk aversion models from economics. Next, we derive a useful property of certainty equivalence functions and it shows how such functions can be naturally applied to sequential auction games. We study the way in which the perceived optimal bidding strategy computed by a risk averse agent, given her probabilistic model of the future, differs from the optimal strategy of a risk neutral agent. Risk averse agents tend to bid more aggressively throughout the sequence of auctions, in order to cover their sunk costs for the initial items in the sequence. However, if the future sequence of auctions is initially perceived as too risky (given the agent’s initial estimation of future closing prices), the best strategy available to a risk averse agent is simply not to participate at all.

Then, we study experimentally the effect that this decision-theoretic bidding behaviour

of risk averse bidders has on his/her expected profit, for markets in which the competition is formed of “myopic”, local bidders (i.e. bidders that require only one particular good). We show that, as expected, more risk-averse bidders have less of a chance to end up with an incomplete bundle, and hence make a loss. But, on average (i.e. assuming a market with repeated interactions), they make less expected profit. When bundles of two possible items are available, we show that more risk averse bidders may prefer to bid for the more common one (even if it has less absolute value), rather than risk the chance of making a loss. This is rational for them, as it reduces their risk, although it also reduces their average expected profit.

The chapter, while providing some important results regarding the complexity of the sequential bidding problem for risk averse agents, leaves several issues to be answered in further work. An important one is deriving optimal bidding strategies in markets in which several synergy agents (i.e. bidders with complementary valuations) bid against each other, not only against myopic, single-value bidders, such as in this work. New bidding heuristics could be developed for software agents that do not only target raw efficiency, but also allow their owners to select a balance between expected profit and risk, based on their personal preferences. Finally, the role of mechanisms such as decommitment [195] and options [120] in reducing or eliminating the exposure problem that risk-averse agents face is another promising direction for further work, which is further explored in Chapter 5 of this thesis.

Chapter 5

Using Priced Options to Solve the Exposure Problem in Sequential Auctions

5.1 Introduction

The previous Chapter introduced the exposure problem that a bidder with complementary valuations, i.e. synergies, faces when she tries to acquire a bundle of goods sold through sequential auctions. Informally, the problem occurs whenever an agent may buy a single good at a price higher than what it is worth to her, in the hope of obtaining extra value through synergy with another good, which is sold in a later auction. However, if she then fails to buy this other good at a profitable price, she ends up with a loss. Otherwise stated, after acquiring the first good, she is exposed to the risk of a potential loss. In the analysis presented in this Chapter (as in much of the previous Chapter), we will henceforth call such a global bidder a *synergy buyer*.

The exposure problem is well known in auction theory and multi-agent systems research. The usual way to tackle this problem in the mechanism design community is to replace sequential allocation with a one-shot mechanism, such as a combinatorial auction [194]. However, this approach has the disadvantage of typically requiring a central point of authority, which handles all the sales. Moreover, many allocation problems occurring in practice (see also Chapters 4 and A of this thesis) are inherently decentralized and sequential. Possible examples range from items sold on Ebay by different sellers, loads appearing over time in distributed transportation logistics, dynamic resource allocation in hospitals, etc.

In Chapter 4, we have already shown how complex the problem that the synergy bidder faces in such a sequence of auctions can be. In that chapter, we have looked at the

problem from the perspective of modeling the bidding decisions of individual agents, based on their aversion to the perceived future risk. However, we did not modify the sequential auction mechanism itself, to try and somehow reduce the exposure to loss that the agents face. In this chapter, we consider a different approach than converting the mechanism to a one-shot, combinatorial auction, one that preserves the sequential nature of such problems, but still reduces considerably the exposure problem that synergy bidders face. The basic idea of our approach (published as [161, 162]), is to allow each seller to auction *options* for the goods she owns, instead of the goods themselves.

Note that this is a very complex problem, and this chapter provides a decision-theoretic analysis of how priced options can be used to address this problem, as well as a first mathematical model to compute option and exercise prices. This is, to our knowledge, the first decision-theoretic analysis for how options should be priced and used in sequential auctions with complementary valued-bidders (pricing models for options in financial markets are very different, as will be explained shortly). As a caveat, we stress that options should not be seen as a “silver bullet” that completely removes the exposure problem, rather, they are a mechanism that, under some assumptions, removes part of the risk exposure and is preferable to both sides (buyers and sellers), by comparison to a direct sale. In fact, auctions for direct sale of the good (as will become apparent in Section 5.1.3) becomes, in our option model, a particular sub-case.

5.1.1 Options: basic definition

An option can be seen as a contract between the buyer and the seller of a good, subject to the following rules:

- The writer or seller of the option has the *obligation* to sell the good for the *exercise price*, but not the right.
- The holder or buyer of the option has the *right* to buy the good for the *exercise price*, but not the obligation.

Since the buyer gains the right to choose in the future whether or not she wants to buy the good, an option comes with an *option price*, which she has to pay regardless of whether she chooses to exercise the option or not.

Options can thus help a synergy buyer reduce the exposure problem she faces. She still has to pay the option price, but if she fails to complete her desired bundle, then she does not have to pay the exercise price as well and thus she limits her loss. So part of the uncertainty of not winning subsequent auctions is transferred to the seller, who may now miss out on the exercise price if the buyer fails to acquire the desired bundle. At the same time, the seller can also benefit indirectly, from the additional participation in the market by additional synergy buyers, who would have otherwise stayed out, because of the exposure to a potential loss.

5.1.2 Related work

In existing multi-agent literature, to our knowledge, there has been only limited work to study the use of options to address the exposure problem.

The first work to introduce an explicit option-based mechanism for sequential-auction allocation of goods to the MAS community is Juda & Parkes [120, 121]. They create a market design in which global bidders are awarded free (i.e. zero-priced) options, in order to cover their exposure problem and, for this setting, they show that truth-telling is a dominant strategy. In this case model, the exposure problem is entirely solved for the synergy buyers, because they do not even have a possible loss consisting of the option price. Having a dominant bidding strategy for the buyers is a very important property, which is difficult to achieve in a model with priced options.

However, the model of Juda & Parkes does have some limitations. First, there may be cases when the market entry effects are not sufficient to motivate the sellers of items to use options. Because the options are assumed to be offered freely (zero-priced), there may be cases in which sellers do not have a sufficient incentive to offer free options, because of the risk of remaining with their items unsold. The sellers could, however, demand a premium (in the form of the option price) to cover their risk. In such cases, only positively-priced options can provide sufficient incentive for both sides to use the mechanism. Furthermore, this free options model relies on the assumption that sellers are always willing to stay in the market longer than buyers.

Priced options have a long history of research in finance (see [107] for an overview). However, the underlying assumption for all financial option pricing models is their dependence on an underlying asset, which has a current, public value that moves independently of the actions of individual agents (e.g. this motion is assumed to be Brownian for Black-Scholes models). This type of assumption does not hold for the online, sequential auctions setting we consider. In our case, each individual synergy buyer has its own private value for the goods/bundles on offer, and bids accordingly.

Another relevant work that studies the use of options in online auctions is that of Gopal et al [87]. Gopal et al. discuss the benefits of using options to increase the expected revenue of a seller of multiple copies of the same good. They do not consider the use of options to solve the exposure problem of buyers with complementary valuations over a bundle of goods (i.e. the synergy buyers in our model). Furthermore, in [87], it is the seller that fixes both the option price and the exercise price when writing the option, which requires rather restrictive assumptions on the behaviour of the bidders.

Finally, there is a connection between options and levelled commitment mechanisms, first proposed by Sandholm & Lesser [195]. In levelled commitment, both parties have the possibility to decommit (i.e. unilaterally break a contract), against paying a pre-agreed decommitment penalty. However, as [195] show, setting the level of the decommitment penalty can be hard, due to the complex game-theoretic reasoning required. There are situations in which both parties would find it beneficial to decommit but neither does, hoping the other party would do so first, to avoid paying the decommitment penalty. This differs from option

contracts, where the right to exercise the option is paid by one party in advance. In our model, this right is sold through an auction, thus the option price is established through an open market.

5.1.3 Outline and contribution of our approach

The goal of this chapter is to study the use of priced options to solve the exposure problem and to identify the settings in which using priced options benefits both the synergy buyer and the seller.

An option consists out of two prices, so an adjustment needs to be made to the standard auction with bids of a single price. The essence of options, in our model, is that buyers obtain the right to buy the good for a certain exercise price in the future. The value of such an option may be different for different market participants at different times. Throughout this study, in order to make the analysis tractable, we have a fixed exercise price and a flexible option price. The seller determines the exercise price of an option for the good she has for sale and then sells this option through a first price auction. Buyers bid for the right to buy this option, i.e. they bid on the option price.

Note that, in this model, direct auctioning of the items appears as a particular sub-case of the proposed mechanism, assuming free disposal on the part of the buyers. If the seller fixes the future exercise price for the option at zero, then a buyer basically bids for the right to get the item for free. Since such an option is always exercised (assuming free disposal), this is basically equivalent to auctioning the item itself.

Based on the above description, we provide both an analytical and an experimental investigation of the setting. Our analysis of the problem can be characterized as decision-theoretic, meaning both buyer and seller reason with respect to expected future price. Our contribution to the literature can be characterized as twofold:

First, we consider a setting in which n complementary-valued goods (or options for them) are auctioned sequentially, assuming there is only one synergy bidder (the rest of the competition is formed by local bidders desiring only one good). For this setting, we show analytically (under some assumptions), that using priced options can increase the expected profit for both the synergy buyer and the seller, compared to the case when the goods are auctioned directly. Furthermore, we derive the equations that provide minimum and maximum bounds between which the bids of the synergy buyer are expected to fall, in order for both sides to have an incentive to use options.

In the second part of the chapter, we consider market settings in which multiple synergy buyers (global bidders) are active simultaneously, and study it through experimental simulations. In such settings, we show that, while some synergy buyers lose because of the extra competition, other synergy buyers may actually benefit, because sellers are forced to fix exercise prices for options at levels which encourages participation of all buyers.

We note also that, while both parts of the paper look into decision theoretic bidding behaviour, we consider different levels of information about the future available to the synergy

bidder. In the analytical case, the exact order of the auctions is assumed to be known, and we consider a bidder that wants a bundle of all the items to be auctioned. In the experimental part, where the synergy bidder only wants a sub-bundle of the goods from a potentially large sequence, we assume that bidding agents only know the number of future buying opportunities for an item of each type, not their exact order. This is actually more realistic for the application scenarios we consider. For example, when bidding to acquire a part-truck order in transportation logistics, it is more realistic to assume that a carrier can approximate the number of future opportunities to buy a complementary load, but not the exact auction order in which future loads will be offered for auction.

The structure for the rest of this chapter is as follows. Sect. 2 lays the foundation for further analysis by deriving the expected profits of synergy buyers and sellers for both the direct sale, respectively for a sale with options. Sect. 3 provides the analytical results and proofs of the chapter, for a market of sequential auctions with one synergy buyer. Sections 4 and 5 summarize the results from our experimental investigations, while Sect. 6 concludes with a discussion.

5.2 Expected profit for a sequence of n auctions and 1 synergy buyer

Section 5.3 will analytically prove, that options can be profitable to both synergy buyer and seller. In order to do that, this section derives the expected profit functions (which depend on the bids of the synergy buyer) for the synergy buyer and the seller. Throughout this study it is assumed that both sellers and buyers are risk neutral and that they want to maximize their expected utility, respectively - in this case - their expected profit.

5.2.1 Profit with n unique goods without options

This section describes the expected profit of the synergy buyer and the sellers as a function of the synergy buyer's bids for a market with n unique, complementary goods, which are sold without options.

Let G be the set of n goods for sale in a temporal sequence of auctions and $v_{syn}(G_{sub})$ be the valuation the synergy buyer has for $G_{sub} \subseteq G$. In this section, we further assume that $v_{syn}(G) > 0$ and $\forall G_{sub} \subsetneq G, v_{syn}(G_{sub}) = 0$. In other words, to somewhat simplify the theoretical model, we consider a synergy buyer that desires the bundle of all the goods considered in the model ($G_{sub} = G$). This assumption will not be used in the experiments.

The goods $G_1..G_n \in G$ are sold individually through sequential, first-price, sealed-bid auctions. Here we choose the auctions to be first price, as they are more tractable to study using game-theoretic analysis. Furthermore, in a sequential setting with valuation complementarities of the agents, second-price auctions do not have the nice dominant strategies properties, described by Vickrey. Furthermore, in many settings where such a model could

be used in practice, such as request-for-quotes (RFQ) auctions in logistics or supply chains, first-price auctioning is often used.

The time these auctions take place in is $t = 1 \dots n$, such that at time t good $G_t \in G$ is auctioned. The above assumptions mean that if the synergy buyer has failed to obtain G_t , then she cannot achieve a bundle, for which she has a positive valuation. So if G_{t+1} is auctioned with a positive reserve price, then obtaining G_{t+1} will only cost the synergy buyer money. Therefore, if the synergy buyer fails to obtain G_t , then it is rational for her to not place bids in subsequent auctions.

The bids of the synergy buyer are $\vec{B} = (b_1, \dots, b_n)$, where b_t is the bid the synergy buyer will place for good G_t , conditional on having won the previous auctions. Because of the first-price auction format, b_t is also the price the synergy buyer has to pay if she has won the auction.

Throughout this analysis, we assume the competition the synergy buyer faces for each good G_t (sold at time t) is formed by local bidders that only require the good G_t . We further assume that these local bidders are myopic, i.e. the bids placed by the synergy buyer have no effect on their bidding behaviour. Therefore, from the perspective of the synergy buyer, the competition can be modeled as a distribution over the expected closing prices at each time point t , more precisely as a distribution over a value bm_t , which is the maximal bid placed by the competition not counting b_t .

Denote by $F_t(b_t)$ the probability that the synergy buyer wins good G_t with bid b_t - where $F_t(b_t)$ depends on whether b_t can outbid the maximal bid bm_t placed by the competition, excluding b_t . For each good G_t , there exists a strictly positive reserve price of $b_{t,res}$, which is the seller's own valuation for that good. Then bm_t is the highest bid of the local bidders (who only want G_t), if that bid is higher than $b_{t,res}$. Otherwise bm_t equals $b_{t,res}$. To deal with ties, we assume the synergy buyer only wins G_t if $b_t > bm_t$ and not if the bids are equal. Then $F_t(b_t)$ can be defined as follows:

$$F_t(b_t) = Prob(b_t > bm_t) \quad (5.1)$$

The synergy buyer only has a strictly positive valuation for the bundle of goods G , which includes all the goods G_t , sold at times $t = 1..n$. Therefore, in a market without options, the a-priori expected profit π_{syn}^{dir} of the synergy buyer is:

$$E(\pi_{syn}^{dir}) = \left[v_{syn}(G) \prod_{i=1}^n F_i(b_i) \right] + \left[\sum_{j=1}^n (-b_j) \prod_{k=1}^j F_k(b_k) \right] \quad (5.2)$$

The synergy buyer wants to maximize her expected profit. So her optimal bids $\vec{B}^* = (b_1^*, \dots, b_n^*)$ maximize equation 5.2:

$$\vec{B}^* = argmax_{\vec{B}} E(\pi_{syn}^{dir}) \quad (5.3)$$

Next the profit of the sellers are examined. It is assumed that all sellers have their own valuation for the good that they sell and that they set their reserve price of $b_{t,res}$ equal to this private valuation. So when the good is sold for b_t , the seller of G_t has a profit π_t^{dir} of $b_t - b_{t,res}$. As previously shown, the synergy buyer only participates when she has won the previous auctions; otherwise bm_t is the maximal placed bid. The expected profit of the seller of the good G_t sold at time t is:

$$\begin{aligned} E(\pi_t^{dir}) = & (E(bm_t) - b_{t,res})(1 - \prod_{i=1}^{t-1} F_i(b_i)) + (F_t(b_t)(b_t - b_{t,res}) \\ & + (1 - F_t(b_t))(E(bm_t | bm_t \geq b_t) - b_{t,res})) \prod_{i=1}^{t-1} F_i(b_i) \end{aligned} \quad (5.4)$$

Intuitively explained, the equation defines the expected utility over 3 disjoint cases: one in which the optimal bids b_i of the synergy bidder were not sufficient to win all auctions up to time t , in which case the expected profit of the seller is the highest expected bid of the local bidders $E(bm_t)$, minus its own reservation value $b_{t,res}$; the second case in which the synergy bidder wins all previous auctions, including the current one (i.e. the one at time t), in which case the expected profit is this bid minus reservation $b_t - b_{t,res}$, and the third in which the synergy buyer won all previous auctions but fails to win the current one, in which case still the highest bid by the local bidders is taken.

5.2.2 Profit with n unique goods with options

Section 5.2.1 derived the expected profit functions for the synergy buyer and the sellers in a market without options. The next step is to do the same for a market with options. This section has the same setting as the general model with n goods being sold, only now an option on G_t is auctioned at time t . Therefore, all the sellers in the market will sell options for their goods, instead of directly the goods themselves. After the n auctions have taken place, the buyers need to determine whether or not they will exercise their option. It is assumed that an option is only exercised if a buyer has obtained her entire, desired bundle. The local bidders are only interested in G_t , so they will always exercise an option on G_t should they have one. The synergy buyer is only interested in a bundle of all goods, so she will only exercise an option (and pay the corresponding exercise price) if she has options on all the goods required.

The option exists out of a fixed exercise price K_t and the synergy buyer's bids on the option price are $\vec{OP} = (op_1, \dots, op_n)$. The maximal bid without the synergy buyer was bm_t , but now opm_t is the maximal placed option price.

Since the competition only wants one good, they do not benefit from having an option and they will always exercise any option they acquire. Therefore the competition's best policy is to keep bidding the same total price, which is the bid without options minus the exercise price. Thus the distribution of the competition is only shifted horizontally to the

left, by the reduction of the exercise price: $opm_t = bm_t - K_t$. Thus, if the synergy buyer bids the same total price (option + exercise), then she has the same probability of winning the auction in both models. Let $F_t^o(op_t)$ be the probability that op_t wins the auction for the option on G_t . So if $op_t + K_t = b_t$, then $F_t^o(op_t) = F_t^o(b_t - K_t) = F_t(b_t)$.

The synergy buyer's expected profit with options then is:

$$\begin{aligned} E(\pi_{syn}^{op}) &= \left[(v_{syn}(G) - \left[\sum_{h=1}^n K_h \right]) \prod_{i=1}^n F_i^o(op_i) \right. \\ &\quad \left. + \left[\sum_{j=1}^n (-op_j) \prod_{k=1}^j F_k^o(op_k) \right] \right] \end{aligned} \quad (5.5)$$

So her optimal bids $\vec{OP}^* = (op_1^*, \dots, op_n^*)$ maximize the profit equation 5.5:

$$\vec{OP}^* = \operatorname{argmax}_{\vec{OP}^*} E(\pi_{syn}^{op}) \quad (5.6)$$

The main difference for the seller of G_t , is that if the synergy buyer wins, then she only earns $K_t - b_{t,res}$ when the option is exercised. She then gains the exercise price, but loses the value the good has to her, which is the reserve price. And the probability of exercise is the probability that the synergy buyer wins all the other auctions. Therefore, the total expected profit of the seller of good G_t sold at time t is:

$$\begin{aligned} E(\pi_t^{op}) &= (E(opm_t) + K_t - b_{t,res})(1 - \prod_{i=1}^{t-1} F_i^o(op_i)) \\ &\quad + \left(F_t^o(op_t) \left[op_t + (K_t - b_{t,res}) \prod_{h=t+1}^n F_h^o(op_h) \right] \right. \\ &\quad \left. + (1 - F_t^o(op_t))(E(opm_t | opm_t \geq op_t) + K_t - b_{t,res}) \right) \prod_{i=1}^{t-1} F_i^o(op_i) \end{aligned} \quad (5.7)$$

Briefly explained, this equation has the same 3-case structure as Eq. 5.4 above. In two cases: when the synergy buyer loses an auction for one the earlier items in the sequence (before the items sold at time t), or when she wins all the earlier auctions, but not the auction at time t , the expected payoffs are equivalents to the direct auctioning case, although this time expressed slightly differently, based on both the exercise and option price. However in one case, when the synergy buyer acquires all the previous items and the current one (middle line in Eq. 5.7), the payoff is composed of two amounts. The option price op_t will be gained for sure, in this case. However, the difference between the exercise and reserve price $K_t - b_{t,res}$ (which signifies the item actually changes hands) is acquired only if the synergy bidder also wins all the subsequent auctions at times $h = t + 1 \dots n$.

This is an important difference, and it would seem from these equations that the seller has no interest to use options, since in one important case, part of the amount she is about to

receive depends on the outcome of future auctions. The key, however, rests in the observation that the synergy buyer should be willing to bid more in total (i.e. $K_t + op_t$) than in the direct auctioning case. This will be analysed in the next Section.

5.3 When options can benefit both synergy buyer and seller

Section 5.2 resulted in the a-priori, expected profit for the synergy buyer and the sellers as a function of the synergy buyer's bids for a market with and without options. This section uses these functions to determine the difference in profit between the two markets, which is $\pi_{\delta t}$ and $\pi_{\delta syn}$ for the seller of good G_t and the synergy buyer respectively, where:

Definition 3

$$\begin{aligned}\pi_{\delta t} &= \pi_t^{op} - \pi_t^{dir}, \\ \pi_{\delta syn} &= \pi_{syn}^{op} - \pi_{syn}^{dir}\end{aligned}$$

So if $\pi_{\delta t}$ and $\pi_{\delta syn}$ are positive, then both agents are better off with options.

5.3.1 When agents are better off with options

Let \vec{B}^* denote the synergy buyer's optimal bidding policy in a market where goods are sold directly (without options). We assume for the rest of Sect. 5.3 that for $1 \leq t \leq n$, $F_t(b_t^*) > 0$ and $F_t(b_t^*) < 1$. So she may complete her bundle, but may also end up paying for a worthless subset of goods. Thus she faces an exposure problem. For the market with options, we define a benchmark strategy \vec{OP}' for the synergy buyer, so that the two markets can easily be compared.

Definition 4 *The benchmark of the synergy buyer's bids with options $\vec{OP}' = (op'_1, \dots, op'_n)$ is that for $1 \leq t \leq n$:*

$$op'_t = b_t^* - K_t$$

In other words, the benchmark strategy implies that the synergy buyer will bid the same total amount for the good, as if she used her optimal bidding policy in a direct sale market. Clearly this does not have to be her profit-maximizing bid in a market where priced options are used. In fact, it is almost always the case that the synergy buyer will bid a different value in a market with priced options. This deviance from the benchmark is denoted by λ_t :

Definition 5 *Let λ_t denote the deviation in the bid of the synergy buyer on the item G_t sold at time t , in a model with options, with respect to her profit-maximizing bid b_t^* in a model without options. So her bid on an option for G_t will be $op'_t + \lambda_t$.*




Figure 5.1: A possible situation in which options are desirable.

These definitions enable us to rigorously define the bounds within which the use of options (with a given exercise price) are desirable for both the synergy buyer and the seller, for each good in the auction sequence (except the last one, for which there is no uncertainty, so the use of options is indifferent). Fig. 5.1 gives the visual description of a generic setting in which options are beneficial for both sides. It shows the possible bids a synergy buyer can place for an option. First, valid bids have to be bigger than the reserve price Res , for each good in the sequence. The point op' is where the synergy buyer keeps bidding the same total price as in a market without options, c.f. Def. 4.

The deviations, in an option model, from the benchmark bid op' is measured by three levels, all denoted with λ : λ_l is the minimal risk premium the seller requires to benefit from using options, λ_h is the maximal extra amount the synergy buyer is willing to pay for an option and $op^* = op' + \lambda^*$ is the synergy buyer's profit-maximizing bid in an option market. So, if it is rational for the synergy buyer to bid an additional quantity between λ_l and λ_h (as shown in Fig. 5.1), then both she and the seller are better off with options.

In the rest of Sect. 5.3, we derive the analytical expressions which can be used to determine the values for λ_l , λ_h and λ^* and compare them. Before this, however, we describe an important assumption behind the proofs in the remainder of this Section.

Assumption used in deriving the proof

Performing an exhaustive theoretical analysis of the minimum, maximum and optimal bidding levels of the λ -s for all auctions in a sequence would not be tractable, as they all influence each other. Therefore, we simplify our proof structure by focusing only on one of the λ parameters: the one corresponding to the first good. This is possible since, as explained in the introduction, each seller sells one good and is only interested in maximizing the expected profit from that sale. The decision of using options contract or a direct sale has to benefit both the seller and the synergy buyer. The buyer must be incentivised to participate in the auction for that particular good, while the seller is incentivised by an additional bidding activity (i.e. higher bid levels) in order to use options.

The reason why we focus on the first good in the sequence is that, for this good, the buyer's probability of not completing her desired bundle, hence her exposure problem, is the greatest. Our proof structure could be generalized as a recursive procedure: if one shows that options are beneficial to use for the first item in a sequence, given a remaining [non-empty] sequence of auctions, this can be generalized to all remaining sub-sequences, (except perhaps, for the very last item, for which the analysis is trivial, as options cannot bring a benefit by comparison to direct sale).

In order to analytically examine the benefits of deviating from the benchmark strategy op'_1 in the first auction, the proofs in this chapter use the additional assumption that the synergy buyer will use the benchmark strategy from Def. 4 for all remaining goods in the sequence. The use of the benchmark bidding strategy for the remaining items can be seen as giving an “upper bound” for the lower lambda value expected by the seller (i.e. λ_l) and a lower bound for the highest value that can be offered by the buyer (i.e. λ_h). We can see this by examining the effect of this assumption on each of the parties:

- For the *synergy buyer*: Being offered the opportunity to use options also in future auctions can only increase his expected profit from future auctions (since $\lambda^* \geq 0$ and $op^* \geq op'$). Otherwise, the synergy buyer will revert to using his benchmark strategy op' , which brings the same expected profit as the direct sale case. His expected profit is at least as large in the options case as in the direct sale case i.e. $E(\pi_{syn,t \geq 2}^{op}) \geq E(\pi_{syn,t \geq 2}^{dir})$.
- For the *seller of the first item*: Because for each of the following items $op^* \geq op'$, the probability that the agent will get all the future items can only increase, for each of the items in the sequence. Formally: $F_h^o(op_h^*) \geq F_h^o(op_h') = F_h(b_h^*), \forall h = 2..n$. This implies that $\prod_{h=2}^n F_h^o(op_h^*) \geq \prod_{h=2}^n F_h(b_h^*)$, therefore the probability that the option for the first item is exercised can only increase. Therefore, this benchmark case acts as a lower bound for the expected profit of the seller, and as an upper bound on the λ_l .

In future auctions the synergy seller and buyer can use options, but this will not negatively affect the initial decisions, i.e. at the beginning of the auction sequence. Therefore, the lambda values referred to in the equations in the following sections could be formally denoted as λ_l^{as} and λ_h^{as} , where in the general case it holds that $\exists \lambda_l, \lambda_h$ such that $\lambda_l \leq \lambda_l^{as}$ and $\lambda_h \geq \lambda_h^{as}$. To avoid overloading the notation, we still use λ_l and λ_h , but the reader should be aware these refer to the tightest bounds on these lambda values, under the assumption that the benchmark bidding strategy is used in all auctions subsequent to the current one.

When synergy buyer is better off with options

This part of Section 5.3.1 examines for which bids the synergy buyer is better off with options. This is done by determining the maximal amount she is willing to pay for options.

Lemma 5.3.1 *Let $\vec{B}^* = \langle b_t^* \rangle$ for $1 \leq t \leq n$ be the vector of optimal bids of the synergy buyer in the model without options, and $op'_t + \lambda_t$ be the bids in a model with options. Then the expected gain (i.e. difference in expected profit) from using options $E(\pi_{\delta syn})$ can be*

written as:

$$\begin{aligned}
E(\pi_{\delta_{syn}}) &= \left[v_{syn}(G) \left(\prod_{i=1}^n F_i(b_i^* + \lambda_i) - \prod_{i=1}^n F_i(b_i^*) \right) \right] \\
&+ \left[\sum_{j=1}^n K_j \left(\prod_{k=1}^j F_k(b_k^* + \lambda_k) - \prod_{i=1}^n F_i(b_i^* + \lambda_i) \right) \right] \\
&+ \sum_{j=1}^n (-\lambda_j) \prod_{k=1}^j F_k(b_k^* + \lambda_k) \\
&+ \left[\sum_{j=1}^n (-b_j^*) \left(\prod_{k=1}^j F_k(b_k^* + \lambda_k) - \prod_{k=1}^j F_k(b_k^*) \right) \right]
\end{aligned}$$

Proof 4 We compute the different in profit between a model with options and a model without options, using expected profit equations (5.5) and (5.2), as defined in the previous section. In a model without options, the optimal bids of the synergy buyer at each time step t are given by b_t^* . In a model with options, we express the bidding policy as a deviation with respect to the benchmark strategy with options, i.e. $op_t' + \lambda_t$. This gives the difference:

$$\begin{aligned}
E(\pi_{\delta_{syn}}) &= \left[(v_{syn}(G) - \left[\sum_{h=1}^n K_h \right]) \prod_{i=1}^n F_i^o(op_i' + \lambda_i) \right] \\
&+ \left[\sum_{j=1}^n (-op_j' + \lambda_j) \prod_{k=1}^j F_k^o(op_k' + \lambda_k) \right] \\
&- \left[v_{syn}(G) \prod_{i=1}^n F_i(b_i) \right] - \left[\sum_{j=1}^n (-b_j^*) \prod_{k=1}^j F_k(b_k^*) \right]
\end{aligned}$$

We can now replace op_t' with the definition of the benchmark strategy (i.e. same total bid amount, as in the case without options), using the properties: $op_t' = b_t^* - K_t$ and $F_t^o(op_t' + \lambda_t) = F_t(b_t^* + \lambda_t)$. This gives:

$$\begin{aligned}
E(\pi_{\delta_{syn}}) &= \left[(v_{syn}(G) - \left[\sum_{h=1}^n K_h \right]) \prod_{i=1}^n F_i(b_i^* + \lambda_i) \right] \\
&+ \left[\sum_{j=1}^n (-b_j^* + K_j - \lambda_j) \prod_{k=1}^j F_k(b_k^* + \lambda_k) \right] \\
&- \left[v_{syn}(G) \prod_{i=1}^n F_i(b_i) \right] - \left[\sum_{j=1}^n (-b_j^*) \prod_{k=1}^j F_k(b_k^*) \right]
\end{aligned}$$

This formula is now re-grouped, separating the terms $v_{synt}G$, $\sum_{j=1}^n K_j$, $\sum_{j=1}^n (-\lambda_j)$ and $\sum_{j=1}^n (-b_j^*)$, each with its corresponding probabilities to complete the proof the proof:

$$\begin{aligned} E(\pi_{\delta synt}) = & \left[v_{synt}(G) \left(\prod_{i=1}^n F_i(b_i^* + \lambda_i) - \prod_{i=1}^n F_i(b_i^*) \right) \right] \\ & + \left[\sum_{j=1}^n K_j \left(\prod_{k=1}^j F_k(b_k^* + \lambda_k) - \prod_{i=1}^n F_i(b_i^* + \lambda_i) \right) \right] \\ & + \sum_{j=1}^n (-\lambda_j) \prod_{k=1}^j F_k(b_k^* + \lambda_k) \\ & + \left[\sum_{j=1}^n (-b_j^*) \left(\prod_{k=1}^j F_k(b_k^* + \lambda_k) - \prod_{k=1}^j F_k(b_k^*) \right) \right] \end{aligned}$$

To explain intuitively Lemma 5.3.1, the difference in expected profits between the two models is formed of 4 parts (corresponding to the 4 lines). First, in an options model, the synergy bidder has a higher probability of getting the desired bundle and extract its value, since she bids more in total (line 1). Furthermore, in an options model, the bidder does not have to pay exercise prices unless she acquires all n items in the desired bundle (line 2). On the minus side, on but she does have to pay a set of additional amounts λ (line 3) for all items she bids on until one is lost (line 3) and, for these items, the chance of acquiring them increases slightly, which also increases the chance of lost bids (line 4).

In the following, we turn our attention to providing equations that allow us to deduce the λ parameters that give the synergy buyer an incentive to use options. As previously explained in Sect. 5.3.1 above, we simplify the proof structure by only focusing on the most important option for the synergy buyer: the one on the first good (when bidding for this good, the probability of not completing her entire bundle is the greatest). This is done under the assumption that for the goods in the sequence, we assume the benchmark strategy is used (i.e. $\lambda_t = 0$ for $t > 1$). For the rest of the items in the sequence, the same proof technique can be applied recursively.

Theorem 5.3.2 *Let λ_1 be the deviation in the bidding strategy, compared to the benchmark strategy op_1^* , as defined in Def. 4. If $\lambda_t = 0$ for $1 < t \leq n$, then by definition, $E(\pi_{\delta synt}) > 0$ if $0 \leq \lambda_1 < \lambda_h$. The value of λ_h (corresponding to $E(\pi_{\delta synt}) = 0$) can be solved as the numerical solution to the following equation:*

$$\begin{aligned} F_1(b_1^* + \lambda_h) \lambda_h = & F_1(b_1^* + \lambda_h) \left[\sum_{j=1}^n K_j \left(\prod_{k=2}^j F_k(b_k^*) - \prod_{i=2}^n F_i(b_i^*) \right) \right] \\ & + (F_1(b_1^* + \lambda_h) - F_1(b_1^*)) \left[v_{synt}(G) \prod_{i=2}^n F_i(b_i^*) - \sum_{j=1}^n (b_j^*) \prod_{k=2}^j F_k(b_k^*) \right] \end{aligned}$$

Proof 5 The proof is based on the difference in profit function derived in Lemma 5.3.1, using the assumption that $\lambda_t = 0$ for $1 < t \leq n$. As the expectation function of the synergy bidder is descending in the value of λ , we determine when $E(\pi_{\delta_{syn}}) = 0$.

$$\begin{aligned}
& \left[v_{syn}(G)(F_1(b_1^* + \lambda_h) - F_1(b_1^*)) \prod_{i=2}^n F_i(b_i^*) \right] \\
& + \left[\sum_{j=1}^n K_j(F_1(b_1^* + \lambda_h) \prod_{k=2}^j F_k(b_k^*)) - (F_1(b_1^* + \lambda_h) \prod_{i=2}^n F_i(b_i^*)) \right] \\
& + (-\lambda_h)F_1(b_1^* + \lambda_h) \\
& + \left[\sum_{j=1}^n (-b_j^*)(F_1(b_1^* + \lambda_h) - F_1(b_1^*)) \prod_{k=2}^j F_k(b_k^*) \right] = 0
\end{aligned}$$

Isolating the values of λ_h yields the formula in Th. 5.3.2.

$$\begin{aligned}
F_1(b_1^* + \lambda_h)\lambda_h &= (F_1(b_1^* + \lambda_h) - F_1(b_1^*)) \left[v_{syn}(G) \prod_{i=2}^n F_i(b_i^*) \right] \\
& + F_1(b_1^* + \lambda_h) \left[\sum_{j=1}^n K_j \left(\prod_{k=2}^j F_k(b_k^*) - \prod_{i=2}^n F_i(b_i^*) \right) \right] \\
& + (F_1(b_1^* + \lambda_h) - F_1(b_1^*)) \left[\sum_{j=1}^n (-b_j^*) \prod_{k=2}^j F_k(b_k^*) \right]
\end{aligned}$$

Which give the following equation for determining λ_h :

$$\begin{aligned}
F_1(b_1^* + \lambda_h)\lambda_h &= F_1(b_1^* + \lambda_h) \left[\sum_{j=1}^n K_j \left(\prod_{k=2}^j F_k(b_k^*) - \prod_{i=2}^n F_i(b_i^*) \right) \right] \\
& + (F_1(b_1^* + \lambda_h) - F_1(b_1^*)) \left[v_{syn}(G) \prod_{i=2}^n F_i(b_i^*) - \sum_{j=1}^n (b_j^*) \prod_{k=2}^j F_k(b_k^*) \right]
\end{aligned}$$

When the first seller is better off with options

We now determine the minimum or lower bound λ_l (the level of λ that, according to Def. 5, keeps the seller of G_1 indifferent about options). In order to compare this bid with the λ_h from the previous section, it is again assumed that $\lambda_t = 0$ for $1 < t \leq n$.

Theorem 5.3.3 If without options the synergy buyer bids \vec{B}^* and with options $op'_1 + \lambda_1$ for

G_1 and op'_t for $1 < t \leq n$, then $E(\pi_{\delta_1})$ for the seller of G_1 is:

$$\begin{aligned} E(\pi_{\delta_1}) &= F_1(b_1^*)(\lambda_1 + (b_{1,res} - K_1) \left[1 - \prod_{h=2}^n F_h(b_h^*)\right]) \\ &\quad + (F_1(b_1^* + \lambda_1) - F_1(b_1^*))(b_1^* + \lambda_1 - E(bm_1 | b_1^* + \lambda_1 \geq bm_1 > b_1^*)) \\ &\quad + (b_{1,res} - K_1) \left[1 - \prod_{h=2}^n F_h(b_h^*)\right] \end{aligned}$$

By definition, λ_1 is the lower bound for λ_l that guarantees that the expected profit of the seller $E(\pi_{\delta_1}) > 0$. The value of λ_l can be obtained as the solution to the equation $E(\pi_{\delta_1}) = 0$, which using the equation above gives:

$$\begin{aligned} F_1(b_1^* + \lambda_l)(-\lambda_l) &= F_1(b_1^* + \lambda_l)((b_{1,res} - K_1) \left[1 - \prod_{h=2}^n F_h(b_h^*)\right]) \\ &\quad + (F_1(b_1^* + \lambda_l) - F_1(b_1^*))(b_1^* - E(bm_1 | b_1^* + \lambda_l \geq bm_1 > b_1^*)) \end{aligned}$$

Proof 6 The difference in profit is equation (5.7) minus equation (5.4):

$$\begin{aligned} E(\pi_1^{op}) - E(\pi_1^{dir}) &= \left(F_1^o(op_1) [op_1 + (K_1 - b_{1,res}) \prod_{h=2}^n F_h^o(op_h)] \right. \\ &\quad \left. + (1 - F_1^o(op_1))(E(opm_1 | opm_1 \geq op_1) + K_1 - b_{1,res}) \right) \\ &\quad - \left(F_1(b_1^*)(b_1^* - b_{1,res}) + (1 - F_1(b_1^*))(E(bm_1 | bm_1 \geq b_1^*) - b_{1,res}) \right) \end{aligned}$$

Recall that the the price op_1 bid in an options model can be expressed in terms of the benchmark strategy op'_1 and the deviation λ_1 .

$$\begin{aligned} E(\pi_{\delta_1}) &= F_1^o(op'_1 + \lambda_1)(op'_1 + \lambda_1 + [(K_1 - b_{1,res}) \prod_{h=2}^n F_h^o(op'_h)]) \\ &\quad + (1 - F_1^o(op'_1 + \lambda_1))(E(opm_1 | opm_1 \geq op'_1 + \lambda_1) + K_1 - b_{1,res}) \\ &\quad - F_1(b_1^*)(b_1^* - b_{1,res}) - (1 - F_1(b_1^*))(E(bm_1 | bm_1 \geq b_1^*) - b_{1,res}) \end{aligned}$$

Furthermore, we can make the substitution to replace op'_1 with its definition, as follows: $op_1 = op'_1 + \lambda_1 = b_1^* - K_1 + \lambda_1$ and $F_1^o(op_1) = F_1^o(op'_1 + \lambda_1) = F_1(b_1^* + \lambda_1)$:

$$\begin{aligned} E(\pi_{\delta_1}) &= F_1(b_1^* + \lambda_1)(b_1^* - K_1 + \lambda_1 + [(K_1 - b_{1,res}) \prod_{h=2}^n F_{oh}(op'_h)]) \\ &\quad + (F_1(b_1^* + \lambda_1) - F_1(b_1^*))(-E(bm_1 | b_1^* + \lambda_1 \geq bm_1 > b_1^*) + b_{1,res}) \\ &\quad - F_1(b_1^*)(b_1^* - b_{1,res}) \end{aligned}$$

Split $F_1(b_1^* + \lambda_1)$ into $F_1(b_1^*)$ and $F_1(b_1^* + \lambda_1) - F_1(b_1^*)$ and combine some K_1 and $b_{1,res}$.

$$\begin{aligned} E(\pi_{\delta 1}) &= F_1(b_1^*)(-K_1 + b_{1,res} + \lambda_1 + \left[(K_1 - b_{1,res}) \prod_{h=2}^n F_h^o(op'_h) \right]) \\ &\quad + (F_1(b_1^* + \lambda_1) - F_1(b_1^*))(b_1^* - K_1 + \lambda_1 + \left[(K_1 - b_{1,res}) \prod_{h=2}^n F_h^o(op'_h) \right]) \\ &\quad - E(bm_1 | b_1^* + \lambda_1 \geq bm_1 > b_1^*) + b_{1,res} \end{aligned}$$

Thus:

$$\begin{aligned} E(\pi_{\delta 1}) &= F_1(b_1^*)(\lambda_1 + (b_{1,res} - K_1) \left[1 - \prod_{h=2}^n F_h(b_h^*) \right]) \\ &\quad + (F_1(b_1^* + \lambda_1) - F_1(b_1^*))(b_1^* + \lambda_1 - E(bm_1 | b_1^* + \lambda_1 \geq bm_1 > b_1^*)) \\ &\quad + (b_{1,res} - K_1) \left[1 - \prod_{h=2}^n F_h(b_h^*) \right] \end{aligned}$$

Since, by definition, $E(\pi_{\delta 1}) = 0$ gives the value of λ_1 , this value can be solved via the equation in Th. 5.3.3.

$$\begin{aligned} F_1(b_1^* + \lambda_l)(-\lambda_l) &= F_1(b_1^* + \lambda_l)((b_{1,res} - K_1) \left[1 - \prod_{h=2}^n F_h(b_h^*) \right]) \\ &\quad + (F_1(b_1^* + \lambda_l) - F_1(b_1^*))(b_1^* - E(bm_1 | b_1^* + \lambda_l \geq bm_1 > b_1^*)) \end{aligned}$$

Intuitively, the difference in profit has two parts: the cases where the synergy buyer wins the auction in both markets and the ones where she only wins with options. With the first, the synergy buyer pays more than she used to and with the second, the synergy buyer pays more than the local bidders, who used to win if $\lambda_1 < \lambda_l$. But both cases have the downside for the seller that the synergy buyer may now not exercise her option.

Both agents can be better off with options

The previous parts of Section 5.3.1 give the equations for the cases when the individual agents are better off with options. These results will now be combined to give the formal conditions when they are both better off. This is done by simply stating that the minimum bid the seller of G_1 requires should be below the maximal value the synergy buyer is willing to pay. As shown the beginning of Section 5.3.1, the equations for λ_l and λ_h that are derived in Theorems 5.3.2 and 5.3.3 above are the narrowest possible interval values, under the assumption that all remaining auctions are direct auctions. Therefore, the solutions to the equations in Thm. 5.3.2 and 5.3.3 are two values λ_h^{as} and λ_l^{as} , where $\exists \lambda_l, \lambda_h$ such that $\lambda_l \leq \lambda_l^{as}$ and $\lambda_h \geq \lambda_h^{as}$. We summarize the results in a final theorem:

Theorem 5.3.4 *Under the condition that the optimal decision of the synergy buyer is to bid λ_x additionally for an option on G_1 (where $\lambda_l^{as} < \lambda_x < \lambda_h^{as}$), then both the seller of G_1 and the synergy buyer have a higher expected profit in a market with only options compared to one without options.*

Proof 7 *The proof of this follows from the previous theorems. Say that the synergy buyer bids $op'_1 + \lambda_x$ for the first good in the sequence, where $\lambda_l^{as} < \lambda_x < \lambda_h^{as}$ and op'_t for the other goods. Then the synergy buyer bids more than $op' + \lambda_l^{as} \geq op' + \lambda_l$ (because $\lambda_l \leq \lambda_l^{as}$), so according to Theorem 5.3.3 the seller of G_1 has a higher expected profit with options. Also, the synergy buyer bids between $0 < \lambda_x \leq \lambda_h^{as} \leq \lambda_h$ extra (as $\lambda_h \geq \lambda_h^{as}$), so according to Theorem 5.3.2 she too has a higher expected profit with options with these bids. Therefore \exists a non-empty interval $[\lambda_l, \lambda_h]$ for which both parties prefer using options, rather than a direct sale.*

5.3.2 Synergy buyer's profit-maximizing bid

In the previous Section, we focused our attention on deriving equations for the bounds λ_l and λ_h between which the additional bids of the synergy buyer have to fall in order for both parties to be incentivised to use options. While these bounds were defined in relation to the expected-profit maximizing bid b^* in a model *without* options, we have not said much about the optimal (i.e. expected profit maximizing) bid op^* in a model with options. The reason for this is that deriving this is much more involved than the optimal policy in a model without options. In this Section, we look at the synergy buyer's profit-maximizing bids op^* , but with the added assumption that $F_1(b_1)$ follows a uniform distribution in the range of the possible bids. Actually, we do this by using the same framework introduced in Def. 5 and Fig. 5.1 above. That means, we compute the deviation λ^* between the optimal bid in a model with options and the optimal bid in a model without options, i.e. the difference $\lambda^* = (K_1 + op_1^*) - b_1^*$ (the reason to do this will become apparent in the proof, but, basically, by taking the difference, several terms drop out). Note that in this section, we still apply the above results and assumption regarding bidding the benchmark strategy in future auctions, but to simplify the notation, we still use λ_l and λ_h , instead of λ_l^{as} and λ_h^{as} .

If the profit-maximizing bid $op_1^* > op'_1 + \lambda_l$, then according to Theorem 5.3.3 the seller of G_1 is better off with options. Therefore, it is in the rational interest of the seller to set the exercise price for selling her good such that the expected optimal bid of her buyers, in a model with options, will provide sufficient incentive for the seller to also use options, and thus the following condition holds: $op_1^* > op'_1 + \lambda_l$. Note that in order to use Theorem 5.3.3, the bids for the other goods are fixed at op'_t . First op_1^* and λ_l are derived.

Lemma 5.3.5 *If $F_1(b_1)$ follows a uniform distribution between ua and ub , then $op_1^* + K_1 -$*

$b_1^* = \lambda^*$, where:

$$\lambda^* = \begin{cases} 0.5(K_1(1 - \prod_{i=2}^n F_i(b_i^*)) + \sum_{j=2}^n K_j(\prod_{k=2}^j F_k(b_k^*) - \prod_{i=2}^n F_i(b_i^*))), \\ \text{if } ua \leq E(\pi_{syn,k \geq 2}^{dir}) \leq ub + (ub - ua) \\ 0, \text{ otherwise} \end{cases}$$

Proof 8 With a uniform bid distribution between ua and ub , the probability of winning with bid b_1 has the following shape:

$$F_1(b_1) = \begin{cases} 0, & \text{if } b_1 < ua \\ (b_1 - ua)/(ub - ua) = \alpha(b_1 - ua), & \text{if } ua \leq b_1 \leq ub \\ 1, & \text{if } b_1 > ub \end{cases} \quad (5.8)$$

$$f_1(b_1) = \begin{cases} 1/(ub - ua) = \alpha, & \text{if } ua \leq b_1 \leq ub \\ 0, & \text{otherwise} \end{cases} \quad (5.9)$$

For F_1^o the variables α_o , ua_o and ub_o are used, where $ua_o = ua - K_1$ and $ub_o = ub - K_1$, so that $F_1(b_1) = F_1^o(op_1)$ when $b_1 - K_1 = op_1$.

First, we determine, for this type of distribution, the equation for the optimal bid b_1^* in a model without options. To do this, we start from the expected profit equation (5.2):

$$E(\pi_{syn}^{dir}) = F_1(b_1)[v_{syn}(G) \prod_{i=2}^n F_i(b_i)] + F_1(b_1)(-b_1) + F_1(b_1)[\sum_{j=2}^n (-b_j) \prod_{k=2}^j F_k(b_k)]$$

$$E(\pi_{syn}^{dir}) = F_1(b_1) \left[-b_1 + [v_{syn}(G) \prod_{i=2}^n F_i(b_i)] + [\sum_{j=2}^n (-b_j) \prod_{k=2}^j F_k(b_k)] \right]$$

So the derivative wrt. b_1 :

$$\begin{aligned} \frac{\partial E(\pi_{syn}^{dir})}{\partial b_1} &= f_1(b_1) \left[-b_1 + [v_{syn}(G) \prod_{i=2}^n F_i(b_i)] \right. \\ &\quad \left. + [\sum_{j=2}^n (-b_j) \prod_{k=2}^j F_k(b_k)] \right] + F_1(b_1)(-1) = 0 \end{aligned}$$

Filling in the equations for f_1 and F_1 leads to:

$$[v_{syn}(G) \prod_{i=2}^n F_i(b_i)] + [\sum_{j=2}^n (-b_j) \prod_{k=2}^j F_k(b_k)] + ua = 2b_1^*$$

Nevertheless, the b_1^* obtained through this formula still has to satisfy the interval constraints $ua \leq b_1^* \leq ub$. This means:

$$ua \leq \frac{[v_{syn}(G) \prod_{i=2}^n F_i(b_i)]}{2} + \frac{[\sum_{j=2}^n (-b_j) \prod_{k=2}^j F_k(b_k)]}{2} + \frac{ua}{2} \leq ub$$

Which yields:

$$ua \leq [v_{syn}(G) \prod_{i=2}^n F_i(b_i)] + [\sum_{j=2}^n (-b_j) \prod_{k=2}^j F_k(b_k)] \leq 2ub - ua$$

Note that the middle expression is, in fact, the expression for the expected profit of a direct synergy bidder, from the second auction onwards (i.e. for $k \geq 2$), discounting the bid to be paid for the first item. Therefore, we can rewrite this condition as:

$$ua \leq E(\pi_{syn,k \geq 2}^{dir}) \leq ub + (ub - ua)$$

From this form, it is easier to explain why outside this interval, $\lambda^* = 0$. If the expected profit of the future sequence $E(\pi_{syn,k \geq 2}^{dir}) < ua$, there is no point in the buyer to continue bidding (either direct or with options), as she cannot afford her desired bundle anyway. Therefore, both b^* and λ^* should be zero. If the expected profit of the future sequence exceeds the value of ub with a whole interval $ub - ua$ (i.e. $E(\pi_{syn,k \geq 2}^{dir}) > ub + (ub - ua)$), then the direct bid assures the bidder of winning the item (as uniform distributions are bounded). But this means that options are also not useful, so again $\lambda^* = 0$ (there is no point of bidding more than in a direct model).

To get the value of λ^* outside these trivial cases is more involved. First, we compute the optimal bid op_1^* in a model with options:

$$E(\pi_{syn}^{op}) = \left[(v_{syn}(G) - \left[\sum_{h=1}^n K_h \right] \prod_{i=1}^n F_i^o(op_i) \right] + \left[\sum_{j=1}^n (-op_j) \prod_{k=1}^j F_k^o(op_k) \right]$$

First, we isolate op_1 in the above equation:

$$\begin{aligned} E(\pi_{syn}^{op}) &= F_1^o(op_1) \left[(v_{syn}(G) - \left[\sum_{h=1}^n K_h \right] \prod_{i=2}^n F_i^o(op_i) \right] \\ &+ F_1^o(op_1)(-op_1) + \left[\sum_{j=2}^n (-op_j) \prod_{k=2}^j F_k^o(op_k) \right] \\ E(\pi_{syn}^{op}) &= F_{o1}(op_1) \left[-op_1 + \left[(v_{syn}(G) - \left[\sum_{h=1}^n K_h \right] \prod_{i=2}^n F_{oi}(op_i) \right] \right. \\ &\left. + \left[\sum_{j=2}^n (-op_j) \prod_{k=2}^j F_{ok}(op_k) \right] \right] \end{aligned}$$

We take the derivative wrt. op_1 :

$$\begin{aligned} \frac{\partial E(\pi_{syn}^{op})}{\partial op_1} &= f_1^o(op_1) \left[-op_1 + [(v_{syn}(G) - [\sum_{h=1}^n K_h]) \prod_{i=2}^n F_i^o(op_i)] \right. \\ &\quad \left. + [\sum_{j=2}^n (-op_j) \prod_{k=2}^j F_k^o(op_k)] \right] + F_1^o(op_1)(-1) = 0 \end{aligned}$$

In order to determine the optimal value op_1^* , we add the condition $\frac{\partial E(\pi_{syn}^{op})}{\partial op_1} = 0$:

$$\begin{aligned} \alpha_o \left[-op_1^* + [(v_{syn}(G) - [\sum_{h=1}^n K_h]) \prod_{i=2}^n F_i^o(op_i)] \right. \\ \left. + [\sum_{j=2}^n (-op_j) \prod_{k=2}^j F_k^o(op_k)] \right] + \alpha_o(op_1^* - ua_o)(-1) = 0 \end{aligned}$$

Which finally yields the following equation for determining op_1^* :

$$[(v_{syn}(G) - \sum_{h=1}^n K_h) \prod_{i=2}^n F_i^o(op_i)] + [\sum_{j=2}^n (-op_j) \prod_{k=2}^j F_k^o(op_k)] + ua_o = 2op_1^*$$

We now focus our attention at computing the difference λ^* between the optima decision-theoretic bid in a model with options vs. a model without options. By definition, we have that: $\lambda^* = (K_1 + op_1^*) - b_1^*$, so $2\lambda^* = 2op_1^* + 2K_1 - 2b_1^*$. When taking this difference, $ua_o = ua - K_1$ and op_k are replaced according to $op_k = op_k' = b_t^* - K_t$ (because for the other auctions, the benchmark strategy is used) and $F_k^o(op_k') = F_1(b_1^*)$. Then all variables cancel each other out, except for the K_t :

$$\begin{aligned} 2(b_1^* + \lambda^* - K_1) &= \left[[(v_{syn}(G) - [\sum_{h=1}^n K_h]) \prod_{i=2}^n F_i(b_i^*)] \right. \\ &\quad \left. + [\sum_{j=2}^n (-b_j^* + K_j) \prod_{k=2}^j F_k(b_k^*)] \right] + ua - K_1 \end{aligned}$$

hence

$$\begin{aligned} 2\lambda^* &= \left[[(v_{syn}(G) - [\sum_{h=1}^n K_h]) \prod_{i=2}^n F_i(b_i^*)] \right. \\ &\quad \left. + [\sum_{j=2}^n (-b_j^* + K_j) \prod_{k=2}^j F_k(b_k^*)] \right] + ua + K_1 - 2b_1^* \end{aligned}$$

thus

$$\begin{aligned} \lambda^* = & 0.5 \left(\left[(v_{syn}(G) - \left[\sum_{h=1}^n K_h \right]) \prod_{i=2}^n F_i(b_i^*) \right] \right. \\ & + \left[\sum_{j=2}^n (-b_j^* + K_j) \prod_{k=2}^j F_k(b_k^*) \right] + ua + K_1 \\ & \left. - \left(\left[v_{syn}(G) \prod_{i=2}^n F_i(b_i) \right] + \left[\sum_{j=2}^n (-b_j) \prod_{k=2}^j F_k(b_k) \right] + ua \right) \right) \end{aligned}$$

After some re-writing:

$$\lambda^* = 0.5 \left(\left(- \sum_{h=1}^n K_h \right) \prod_{i=2}^n F_i(b_i^*) + \sum_{j=2}^n K_j \prod_{k=2}^j F_k(b_k^*) + K_1 \right)$$

Re-arranging the parantheses:

$$\lambda^* = 0.5 \left(K_1 - K_1 \prod_{i=2}^n F_i(b_i^*) - \sum_{h=2}^n K_h \prod_{i=2}^n F_i(b_i^*) + \sum_{j=2}^n K_j \prod_{k=2}^j F_k(b_k^*) \right)$$

Which finally leads to the equation in Lemma 5.3.5:

$$\lambda^* = 0.5 \left(K_1 \left(1 - \prod_{i=2}^n F_i(b_i^*) \right) + \sum_{j=2}^n K_j \left(\prod_{k=2}^j F_k(b_k^*) - \prod_{i=2}^n F_i(b_i^*) \right) \right) \quad (5.10)$$

The main intuition behind this formula is that, in an options model, the synergy buyer saves the exercise price when she fails to complete her bundle. Therefore, it is her profit-optimizing strategy, in a model with options, to increase her bid with a part of the potential savings on the exercise prices of subsequent auctions.

Lemma 5.3.6 *If $F_1(b_1)$ follows a uniform distribution, then the lower bound is:*

$$\begin{aligned} \lambda_l = & -(b_1^* - ua + \left[1 - \prod_{h=2}^n F_h(b_h^*) \right] (b_{1,res} - K_1)) + \\ & + \sqrt{(b_1^* - ua + \left[1 - \prod_{h=2}^n F_h(b_h^*) \right] (b_{1,res} - K_1))^2} \\ & - 2(b_1^* - ua) \left[1 - \prod_{h=2}^n F_h(b_h^*) \right] (b_{1,res} - K_1) \end{aligned}$$

Proof 9 Take the λ_l equation from Theorem 5.3.3. With a uniform distribution, $F_1(b_1) = \alpha(b_1^* - ua)$ and $E(bm_1|b_1^* + \lambda_l \geq bm_1 > b_1^*) = b_1^* + 0.5\lambda_l$. So the equation becomes:

$$\begin{aligned} \alpha(b_1^* + \lambda_l - ua)(-\lambda_l) &= \alpha(b_1^* + \lambda_l - ua)((b_{1,res} - K_1) \left[1 - \prod_{h=2}^n F_h(b_h^*)\right]) \\ &+ \alpha\lambda_l(b_1^* - b_1^* - 0.5\lambda_l) \end{aligned}$$

Dividing both sides by α and reducing b_1^* in the last parenthesis gives:

$$(b_1^* + \lambda_l - ua)(-\lambda_l) = (b_1^* + \lambda_l - ua)((b_{1,res} - K_1) \left[1 - \prod_{h=2}^n F_h(b_h^*)\right]) + \lambda_l(-0.5\lambda_l)$$

After re-arranging the terms and moving the left-hand side to the right, this yields:

$$(b_1^* + \lambda_l - ua)(\lambda_l + (b_{1,res} - K_1) \left[1 - \prod_{h=2}^n F_h(b_h^*)\right]) - 0.5\lambda_l^2 = 0$$

The above equation can be brought to standard, 2nd order polynomial form in the unknown λ_l :

$$\begin{aligned} 0 &= 0.5\lambda_l^2 + \lambda_l(b_1^* - ua + (b_{1,res} - K_1) \left[1 - \prod_{h=2}^n F_h(b_h^*)\right]) \\ &+ (b_1^* - ua)((b_{1,res} - K_1) \left[1 - \prod_{h=2}^n F_h(b_h^*)\right]) \end{aligned}$$

This polynomial equation can then be solved via the quadratic formula:

$$\begin{aligned} \lambda_l &= -(b_1^* - ua + \left[1 - \prod_{h=2}^n F_h(b_h^*)\right](b_{1,res} - K_1)) \\ &\pm \sqrt{(b_1^* - ua + \left[1 - \prod_{h=2}^n F_h(b_h^*)\right](b_{1,res} - K_1))^2 - 2(b_1^* - ua) \left[1 - \prod_{h=2}^n F_h(b_h^*)\right](b_{1,res} - K_1)} \end{aligned}$$

Note that formally, we the condition $ua \leq b^* \leq ub$ should also be imposed in the above equation. However, if b^* falls for the direct sale case falls outside this interval (i.e. if $\leq E(\pi_{syn,k \geq 2}^{dir}) < ua$ or $\leq E(\pi_{syn,k \geq 2}^{dir}) > ub + (ub - ua)$), we know that the the lambda of the seller $\lambda^* = 0$, so there is no point in the seller even considering offering options. Outside this interval, it makes no sense to compute an expression for λ_l .

The next and final step in these proofs should actually involve comparing the equations for λ^* (from Lemma 5.3.5) and λ_l (from Lemma 5.3.6), such as to derive a condition for when $\lambda_l < \lambda^*$. We found that getting a closed form expression for this condition is not possible for these two equations. However, the framework developed above is sufficient to enable the seller to solve this condition numerically using a standard solver and, thus, choose the optimal level for the exercise price K_1 .

5.4 Simulation of a market with a single synergy buyer

This section presents an experimental examination of a market with one synergy buyer. It introduces the market entry effects in the synergy buyer's behaviour, as well as the threshold effects that may determine which exercise prices the seller chooses for her options. This experimental analysis is performed here for a market with one synergy bidder and several local bidders, while Sect. 5.6 considers a market with multiple synergy bidders.

The experimental setting is as follows: we consider a simulation where two goods A and B are auctioned n_A and n_B times respectively. The synergy buyer desires one copy of both goods and has zero valuation for the individual goods. That is, each synergy (or global) bidder requires exactly one bundle of $\{A, B\}$. Note this is somewhat different than the setting used in Chapter 4, in which a bundle of the same good was needed.¹ In the setting considered in this Section, local bidders only want one good and participate in one auction, thus their bids can be modeled as a distribution.

Furthermore, in order to simplify the analysis of the model, we assume there is a single seller who auctions all the goods. This is actually equivalent to studying whether *on average* sellers have an incentive to use options. To explain, on any single sequence of auctions taken in isolation, the sellers of different items may have highly diverging incentives to use options, based on their position in the auction queue. However, in a very large setting, where buyers enter the market randomly, it is difficult for any individual seller to strategize about her particular place in the sequence (and, furthermore, in most markets she may simply have no information to do this). Our goal is to study under which conditions, on average, sellers benefit from using options if there are synergy buyers in the market. Also, to somewhat reduce the number of test parameters, we further assume that the exercise price is the same for all goods of the same type. So the seller needs to determine which exercise price for A and which for B maximize her expected profit.

Note that, typically a seller has a resale value of for the goods that remain unsold, which is typically lower than the value at the start of the auction sequence. The reason for this may be that there is some time discounting associated with waiting for a sequence of auctions to resell her items, or even a listing cost, which is paid per auction (such as in the Ebay case). In this paper, we do not explicitly simulate resale, but we use a reservation value, which represents the expected resale value the seller expects to get, if she is forced to resell her items.

To summarize, simulations were run in Matlab and had the following parameters:

¹An intuitive way to think about this setting is as a sequential sale of individual shoes of exactly the same type, where A is the left shoe, and B is the right shoe, and each synergy buyer requires exactly one pair.

Name	Explanation
n	The number of auctions.
$mean$	The mean of price distribution.
std	The standard deviation of price distribution.
res	Reserve prices.
v_{syn}	Valuation of synergy buyer for A and B combined.
k	Number of simulations for each auction run (i.e. how many times a sequence of auctions is repeated for one set of parameters).

A basic simulation run is as follows. First, all possible auction sequences are determined for the given number of auctions for A and B. The simulation is then run for all these sequences, both for a direct sale setting and for a setting where the items are sold through options with given exercise prices.

For each auction, in each simulation run, there is a set of local bidders, assumed myopic. The bids of these local bidders are therefore, assumed to follow a normal price distribution, with the parameters n , $mean$, std and res consisting out of two values: one for good A and one for good B. For each simulation run, the synergy bidders(s) are asked to determine their profit-maximizing bid for that setting, as described in the next section. The optimization required for determining their optimal bid is done using the Matlab function “fminsearch” from the Optimization Toolbox.

Since there may be considerable variance in the bids of the local bidders (which are myopic) each possible auction sequence is run k times (typically, we had $k > 10000$). The average profit of the seller and the synergy buyer which are reported here, for both the case of with and without options, are averages over all these k simulations and also over all possible auction orders of items A and B in the sequence.

5.4.1 Synergy buyer’s bid strategy

This section describes how the synergy buyer determines her bids in the simulation. In order to neutralize the effect that the exact order items are auctioned in plays on the bidding strategy, we add the assumption that the synergy buyer knows the number of remaining auctions, but not the order they will be held in. This remaining number of auctions of each type is common knowledge (i.e. the synergy bidders can always observe how many auctions of each type are left before they have to leave the market, and so does the seller).

The model described here is for a situation without options. But in order to apply it to a situation with options, one merely has to replace the variables: $b_t = op_t - K_t$ and $v_{syn}(A, B) := v_{syn}(A, B) - K_A - K_B$. As in the analytical section, we assume a bidder only wants a complete bundle of $\{A, B\}$. Therefore, $v_{syn}(A) = 0 = v_{syn}(B)$.

Determining the synergy buyer’s profit-maximizing bid b_t^* at state t basically involves solving the Markov Decision Process (MDP), where we select the optimal bid b_t^* at time t , subject to the optimal bid b_{t+1}^* being selected for the future time point $t + 1$ (which in

this case, is an auction). We can, however, use the valuation function of the bidding agent to significantly reduce the state space of the MDP, as shown below. However, first we introduce some notation.

Let b^* be the immediate best response to the state, which depends on four variables: z_A, z_B, X and I_t . The variables z_A and z_B are the number of remaining auctions for A and B respectively (including the current auction), so $z_A \leq n_A, z_B \leq n_B$. The type of good, which is currently sold, is denoted by I_t . The set of goods the synergy buyer owns (i.e. the endowment) is described by X , which can either be $\emptyset, \{A\}$ or $\{B\}$. If X is $\{A, B\}$ then the synergy buyer is done. Let $Q(z_A, z_B, X, I_t, b_t)$ be the expected profit of the synergy buyer when bidding b_t . Note that, in these definitions, b_{t+1}^* and $V_{t+1}()$ denote the best available bid, respectively best expected value for the next state (as computed by recursion), while I_{t+1} is the type of the next item in the auction sequence. Therefore, using MDP notation, the profit-maximizing bid b_t^* is determined as follows:

$$b_t^* = \operatorname{argmax}_{b_t} Q(z_A, z_B, X, I_t, b_t) \quad (5.11)$$

Where the expected profit is determined via:

$$\begin{aligned} Q(z_A, z_B, X, I_t = A, b_{t+1}^*) &= F_A(b_t)(-b_t \\ &+ V_{t+1}(z_A - 1, z_B, X \cup A, b_{t+1}^*)) + (1 - F_A(b_t))V_{t+1}(z_A - 1, z_B, X, b_{t+1}^*) \end{aligned} \quad (5.12)$$

$$\begin{aligned} Q(z_A, z_B, X, I_t = B, b_t) &= F_B(b_t)(-b_t \\ &+ V_{t+1}(z_A, z_B - 1, X \cup B, b_{t+1}^*)) + (1 - F_B(b_t))V_{t+1}(z_A, z_B - 1, X, b_{t+1}^*) \end{aligned} \quad (5.13)$$

Where $V()$ is the value of a state, which simply means the maximum expected profit of that state:

$$V_t(z_A, z_B, X, b_t) = \max_{b_t} Q(z_A, z_B, X, I_t, b_t) \quad (5.14)$$

Looking at the formula for $Q()$, it basically says that for the probability of winning the auction with her bid, the synergy buyer has to pay a price equal to her bid and the good is included in the endowment X of the next state. If she does not win the auction, then the value of the current state is equal to the value of the next state.

As we mentioned before, in computing its optimal bidding strategy used in the experimental Section, we assume the synergy buyer does not know whether the next auction will be for A or B , she only knows the total numbers of auctions for A and B remaining. We acknowledge this is a departure from the formulas in the theoretical analysis, where the exact order of the auctions was taken into account to compute the bidding strategies. There are two reasons to use this assumption here. The first is that it reduces considerably the state space that needs to be modeled when computed the optimization. But the second is that we also find this choice more realistic if this model is to be applied to real-life settings. For example, when bidding on a part-truck order in a logistic scenario, it is more realistic to assume that

a carrier can approximate the number of future opportunities to buy a complementary load, but not the exact auction order in which future loads will be offered for auction.

If we assume the synergy buyer only knows the total numbers of auctions for A and B remaining (and not their exact order), then her bidding strategy is based on assuming each future auction has an equal probability to occur. Therefore, the probability of an auction for A occurring next is simply the number of remaining auctions A divided by the total number of remaining auctions. Thus, a weighted average can be used to determine the value of the next auction, while not knowing for which good it will be for.

Apart from this general framework, we can prune the state space with the cases in which we know the synergy buyer's bid is zero:

$$b_t^* = \operatorname{argmax}_{b_t} Q(0, z_B, X, B, b_t) = 0, \text{ with } A \notin X \quad (5.15)$$

$$b_t^* = \operatorname{argmax}_{b_t} Q(z_A, 0, X, A, b_t) = 0, \text{ with } B \notin X \quad (5.16)$$

$$b_t^* = \operatorname{argmax}_{b_t} Q(z_A, z_B, X, I_t \in X, b_t) = 0 \quad (5.17)$$

With the first two cases, the synergy buyer can no longer obtain her desired bundle, because she does not own the complementary item and there is no chance left of acquiring it. The last equation is for the case when the synergy buyer already has a copy of the type of good (and, from her valuation function, she only wants exactly one copy of A and B). The corresponding values of these states are:

$$V(0, z_B, X, b_t^*) = 0, \text{ if } A \notin X \quad (5.18)$$

$$V(z_A, 0, X, b_t^*) = 0, \text{ if } B \notin X \quad (5.19)$$

$$V(z_A, z_B, \{A\}, b_t^*) = V(0, z_B, \{A\}, b_t^*) \quad (5.20)$$

$$V(z_A, z_B, \{B\}, b_t^*) = V(z_A, 0, \{B\}, b_t^*) \quad (5.21)$$

The first two equations correspond to the case when the buyer can no longer get the complementary-valued item, therefore the sequence of auctions of the same type has no value to her. In both these cases $b_t^* = 0$. The last two equations are important, since they help the most to reduce the state space. Basically, as already mentioned, we assume that a synergy bidder only wants exactly one bundle of $\{A, B\}$. If she already owns a good of one of the two types, she will no longer be interested in the remaining auctions for that type of good. Therefore, the valuation $V()$ of these states is equivalent to a state when no auctions are remaining for the type of good she already owns (as she would not take part in those anyway). All these techniques help reduce the recursive search.

To conclude, to determine the synergy buyer's bids in any situation, the values of b_t^* and $V()$ need to be calculated for the following states:

$$\forall z_B > 0 \quad Q(0, z_B, \{A\}, B, b_t)$$

$$\forall z_A > 0 \quad Q(z_A, 0, \{B\}, A, b_t)$$

$$\forall z_A > 0, z_B > 0 \quad Q(z_A, z_B, \emptyset, A, b_t)$$

$$\forall z_A > 0, z_B > 0 \quad Q(z_A, z_B, \emptyset, B, b_t)$$

5.4.2 Experimental results: market entry effect for one synergy buyer

First, we study experimentally the incentives to use options for the sellers and buyers, in the case there is just one synergy bidder present in the market. In order to study different dimensions of such markets, we considered several combinations of parameter settings.

The first setting has $n_A = 2$ and $n_B = 2$. As mentioned above, the local bidders are considered myopic and only bid in one local auction. Therefore, their bids can be modeled as a distribution $\sim N(10, 4)$ for both goods. The goods A and B are, in this model, of equal rarity and attract an equal amount of independent competition during bidding. This choice is not random, as having a certain degree of symmetry in the experimental model allows us to reduce the number of parameter settings we need to consider. More specifically, we assume the same exercise prices are set for both goods of type A and B. This is a reasonable assumption, because A and B are of symmetric value and because bidders do not know in advance the exact order goods will be sold in.

Furthermore, for each good, the seller has a reservation value $res = 8$, which gives its estimate resell value in the case the synergy buyer acquires an option for the item, but fails to exercise it. Since, on average, myopic bidders bid have an expected mean of 10 for an item, 20% is a reasonably safe estimate of a resell value.

The value of a bundle of $\{A, B\}$ for the synergy buyer is an important choice, especially in relation to the mean expectation μ of the bids placed by single-item bidders. We considered two settings: $v(A, B) = 24$ (thus 20% more, on average, than local competition) - with results shown in Fig. 5.2, and $v(A, B) = 21$ (which is only 5% more on average than local competition) - with results shown in Fig. 5.3.

Looking at these two figures, some important effect can be observed. First, we mention that the seller has an immediately higher expected profit with options compared to direct sale. This is because an option is sometimes not exercised and then the seller gets to keep the good (for which she has a positive valuation), while the synergy buyer still pays the option price.

There are two main effects to be observed from Fig. 5.2 and 5.3:

- First, the synergy buyer in such a market always prefers *higher* exercise prices (an effect clearly seen in both Figs. 5.2 and 5.3). This may be counter-intuitive at first, but is a rational expectation. If the option for an item is sold with a higher exercise price, then the synergy buyer can bid more aggressively on the option price to get the item, since she is “covered” for the loss represented by the exercise price. The myopic bidders extract no advantage from being offered the good as an options vs. a direct sale, because, if they acquire the option, they would always exercise it regardless. Therefore, they will simply lower their bid for the option with the amount represented by the exercise price.
- Second, the expected profit of the seller seems to decrease between intervals if she has to sell the option with a higher exercise price. The main reason for this is that there is



Figure 5.2: Percentage increase in profit for a model using options wrt. direct sale, for the case there is one synergy buyer is present in the market. In the setting, there are two items of type A sold and two items of type B. For all 4 items, the bids of the local bidders follow the distribution $N(10, 4)$, while the valuation of the synergy buyer is $v(A, B) = 24$ (thus 20% more, on average, than the local bidders). What is varied on the horizontal axis is the exercise price with which the items are sold (assuming they are set the same for all items, being of equal rarity). Note that the figure is super-imposed: the left-hand side axis refers exclusively to the seller, while the right-hand side axis refers exclusively to the synergy bidder. From this picture, one can already see the important effect: synergy buyer prefers, on average, higher exercise prices, while seller prefers lower ones. Note that there is a sudden increase in profit, on the seller side, for the options case with $k = \epsilon > 0$, wrt. direct auctioning. This is simply because, with options, the seller gets to keep the item (for which it has a non-residual value), rather than the buyer, who disposes of it (as in the direct sale case).

some chance that she or she would remain with her item unsold (because the option is not exercised), and thus only extract her reservation value for that item. There is, however, an important difference between the cases shown in Fig. 5.2 and 5.3, which is the participation thresholds (that appear as “peaks” in the picture), where the expected profit of the seller seems to “jump” at a new level. These can be explained by the synergy buyer joining the market, as the expected profit becomes non-negative. The threshold nature is determined by the discrete nature of the auction sequence, as is explained below.

Such a participation threshold is illustrated in Fig. 5.3 is the increase in the seller’s expected profit when the exercise price is set above a certain level ($K \geq 2.5$, for the settings in Fig. 5.3). Such thresholds can be explained as follows. If the synergy buyer currently



Figure 5.3: Percentage increase in profit for a model using options wrt. direct sale, for the case there is one synergy buyer is present in the market. The settings are exactly the same as those is in Fig. 5.2 above: 2 auctions for A and 2 for B, with local, myopic bidders following $N(10, 4)$. However, now the valuation of the synergy buyer is $v(A, B) = 21$ (thus only 5% more, on average, than the local bidders). One can see, however, that there is an important difference by comparison to Fig. 5.2: the threshold effect in the profit increase for the seller when the exercise price $K \geq 2.5$. Intuitively, the reason this effect occurs is the market-entry effect on the part of the synergy buyer, who would otherwise stay out for this lower valuation

owns nothing, then she will only bid on a good if the number of remaining auctions and their exercise prices give her a prior expectation of a positive profit. Conversely, if the synergy buyer is not offered a sequence of option sales from which she derives a positive expected profit, she has the incentive to leave the market altogether. There are two main factors that increase a synergy buyer's expected profit in a sequence of auctions (sold as options):

- The number of remaining future auctions of the other good, necessary to complete her bundle.
- The exercise price of the options (that only needs to be paid at the end). This should be high enough to cover the risk, given her valuation for the bundle.

Note that in some market setting (such as the one in Fig. 5.2), no participation effects (i.e.e. thresholds) occur, because the value the synergy buyer assigns to her desired bundle is already high enough, so she would participate in the market anyway (i.e. regardless of whether she gets offered options or not), and at any point in the sequence that there is still a chance of completing her bundle.

However, in the valuation settings in Fig. 5.3, the synergy buyer will only bid on a good




Figure 5.4: Percentage increase in profit for the case of one synergy buyer, for longer auction sequences. The settings in terms of valuations are exactly the same as those in Fig. 5.3 above: the synergy buyer has a value $v(A, B) = 21$, while single-item bidders bid according to $N(10, 4)$. One change is that now there are 4 auctions available for each type, i.e. 4 auctions for an item of type A and 4 for B. Notice that now there are multiple thresholds, since there are multiple points when the market entry effect of the synergy buyers appears. However, on average, the percentage increases in expected profits for the synergy buyers are lower, when compared to the direct auctions case. The reason for this is that, with multiple future buying opportunities, the exposure problems that synergy bidder faces decreases.

if there are two remaining auctions for the other good. So she places a bid for A if the auctions are $[A, B, B]$, but not if they are $[A, B]$. This is because with a single auction for B, the risk of ending up with a only a worthless A is too great. But in a market with exercise prices of at least 2.5, the risk is reduced and one remaining auction is already enough for the synergy buyer to stay in the market. So a higher exercise price enables the synergy buyer to stay the market, even if she owns nothing and there are only a few auctions left, which increases the seller's expected profit. This increase in participation is beneficial to the seller, who thus has an incentive to fix the exercise prices $K_A = K_B = 2.5$.

5.5 Settings with longer sequences of auctions and effect of auction order

In the previous Section, we examined a sequence of auctions of a specific length of $n_A = 2, n_B = 2$. We now look at whether we can observe similar effects in the case when the number of opportunities to buy goods A and B increases. With the exception of auction

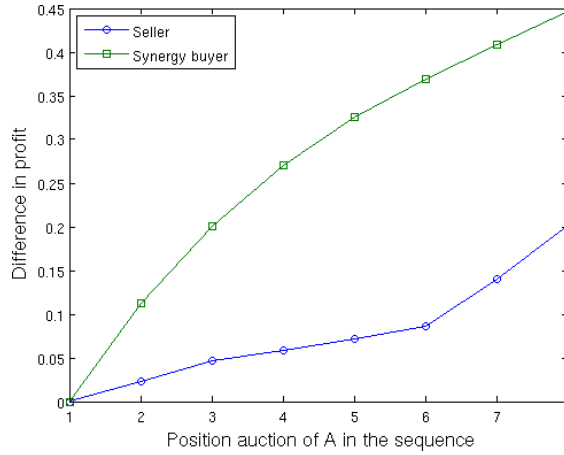


Figure 5.5: Influence of the position in an auction queue of an item on the seller's expected profit. Settings are the same as in Fig. 5.2, but with one important difference: the rarity of the goods is no longer symmetric. There is now only 1 auction for a good of type A, but 7 auctions for a good of type B. What is varied along the horizontal axis is the position in the auction queue of the sale of the rarer item (of type A). The graph shows the absolute difference in profit for a seller of an item of type B and for the synergy buyer (i.e. the difference in profit between an options and direct auctions model). Note that, if the rare item of type A is sold at the end of the auction sequence, the benefit of selling item B through an option increases, because the exposure risk of not acquiring item of type A increases.

lengths, the parameters are kept the same as in the previous case. First, we keep the relative rarity of both goods symmetrical, but increase the number of auctions available for each to 4, i.e. $n_A = n_B = 4$. Results are shown in Fig. 5.4.

Basically, there are two main effects to observe here. First, the benefits to the buyer of having options mechanism decreases (seen from comparing the percentage increases shown in the right-hand vertical axis of Figs. 5.3 and 5.4). The reason for this (as discussed in the earlier, risk-based bidding chapter) is that, in sequential auctions, the number of available future opportunities plays a big role in how big the exposure problem the synergy buyer faces is. If there is less exposure, then the relative benefits of using options becomes smaller (although it is still quite considerable). The second effect to be observed from Fig. 5.4 is that there are more participation thresholds (denoted by peaks), but they are smaller. The reason is that, for a longer sequence of auctions, there are more possible sequences of remaining auction combinations. The synergy bidder will join in the bidding in some, but not in others, leading to multiple participation thresholds.

The second problem we look in this subsection at is what happens if the relative frequency of the two goods is more asymmetric. We keep the same total number of auctions

in the sequence (8), but the relative frequency is highly asymmetric: $n_A = 1, n_B = 7$. As mentioned, in the previous graphs, results were averaged over all possible auction orders - while here, by contrast, we look at auction orders one by one.

For this setting, there are exactly 8 possible auction orders, corresponding to the point where the rarer good (type A) can be inserted in the auction queue. What is varied on the horizontal axis is this position of the type A good. The reason why we look at whether a seller of items of type B would use options is that the exposure of the synergy buyer exists for the other good in the sequence. For the single item of type A, the benefits of using options are limited, because the synergy buyer has 7 other auctions in which to acquire the second item anyway, hence she has much less of an exposure problem.

Clearly, we can see an important effect of the position of the rarer good in the auction queue, from the perspective of both parties. If the item of type A is sold at the very beginning of the auction sequence, then the synergy bidder has no exposure problem left for the rest of the sequence, hence there is no incentive to use options, for either party. However, it is at the very end of the auction sequence, the synergy buyer will not know whether she would need the item acquired until all auctions end. For this case, the benefits of using options are considerably greater.

5.6 Multiple synergy buyers

Finally, we consider market settings in which multiple synergy buyers are active simultaneously. Much of the experimental set-up and parameter choices are the same as described in the above Sections, for the case of one for the single synergy buyer. The only difference is that now multiple synergy buyers may enter and leave the market at different times and they have different valuations for the combination of A and B.

We have to emphasize that the results from this Section are still rather preliminary and are based on some restrictions on the reasoning capability of the synergy buyers in the market. Specifically, as in the single-bidder case, we assume the synergy bidders have some prior expectations about the closing prices in future auctions and compute their optimal strategy wrt. this expectation. In these results, this expectation is assumed the same for all synergy bidders, which is a reasonable choice in comparing their strategies, but assuming the sequence of auctions considered is too short for other synergy buyers to learn about existing competition and adapt their bids. In a more realistic market, however, synergy bidders could be expected to be able to learn and adjust their expectations based on past interactions, as well as reason game-theoretically about the fact that another synergy bidder may present in the market at the same time. At this point, these more sophisticated forms of reasoning are left to future work.

As in the previous section all simulations of this section have reserve prices of 8 and local bidders following $\sim N(10, 2.5)$. The first two experiments also have two synergy buyers syn_1 and syn_2 with valuations for both goods of 21.5 and 22.5 respectively. The




Figure 5.6: Percentage increase in profits for a market with 2 synergy bidders. There are 3 auctions for A and 3 for B, and for each one the bids from the competition formed by local bidders follows the distribution $N(10, 2.5)$. The valuations of the two synergy bidders for a bundle $\{A, B\}$ are 21.1 for *syn1*, respectively 22.5 for *syn2*. The order the agents enter the market is described by Fig. 5.8 below (so the two agents do not compete directly against each other in this setting). Notice that, in this case, the average profit of *syn2* does not decrease with the entry of *syn1* in the market.

order the synergy bidders enter the market (and the number of auctions they can stay in) are given in Figs. 5.8 and 5.8, while results for all settings are shown in Fig. 5.6, respectively 5.7. In the following, we will discuss these in separate subsections.

5.6.1 Two synergy buyers interacting indirectly through the exercise price level

In the setting examined here, the two synergy buyers each have $n_A = 3$ and $n_B = 3$, without the other agent participating in these auctions. An example of such an auction sequence is shown in Fig. 5.8. However, these two synergy bidders do interact indirectly as follows. Since options are sold through open auctions based on the option price, the seller has to fix the exercise prices for the whole market (i.e. for all auctions in the sequence). So while synergy buyers may not participate in the same auctions, their presence does influence the competition through the exercise prices set by the seller.

This effect can be seen in Fig. 5.6, in which the seller maximizes her expected profit at $K = K_A = K_B = 2.4$. In this case *syn2* is better off, because without the presence of *syn1* she would be offered options with lower exercise prices. But *syn1* is worse off, because if she were alone in the market the seller would choose $K = 3.2$, which gives her a higher expected profit. Yet, due to *syn2*, the seller sets $K = 2.4$. In this case, due to the seller's




Figure 5.7: Percentage increase in profits for a market with with 2 synergy bidders. The setting and valuations are the same as in Fig. 5.6 above. However, the order the agents enter the market is now described by Fig. 5.9 below (so the two agents *do* compete directly for the same goods). Notice that, in this case, the average profit of *syn2* decreases due to the additional competition from *syn1*.




Figure 5.8: An auction sequence for the case shown in Fig 5.6.

choice of exercise prices, one synergy buyer (*syn₁*) gains, while *syn₂* loses.

5.6.2 Direct synergy buyer competition in the same market

Next, we considered a setting in which synergy buyers compete directly for some of the goods. The entry points for such a setting are shown in Fig. 5.9, while simulation results are given in Fig. 5.7.




Figure 5.9: An auction sequence for the case shown in Fig. 5.7.

As can be seen in Figure 5.7, the profit of syn_2 drops at 2.5. In previous figures the synergy buyers' profits were monotonically increasing in the exercise prices, because they then have a smaller loss when they fail to complete their bundle. But now this effect cannot immediately compensate the extra competition coming from syn_1 , who participates in the same auctions more often after this threshold at 2.5. So, in this case, both synergy buyers lose from the presence of additional bidders. While one synergy buyer (i.e. syn_2) should benefit because she is offered better (higher) exercise prices than if she were alone in the market, this effect cannot immediately compensate the additional competition.

5.6.3 Larger simulation with random synergy buyers' market entry

In the final results we report in this Chapter, we conducted a larger scale simulation with multiple synergy buyers, which can enter the market randomly, with a certain probability.

The experimental setup implies that each sequence of auctions (forming a test case) has 10 items of each type (i.e. $n_A = 10$ and $n_B = 10$). What differs from previous settings is the random entry of synergy buyers. For each auction, there is a 25% chance that a synergy buyer will enter the market. If she does, then her valuation is drawn from a uniform distribution between 20 and 22 and she will stay in the market for exactly four auctions. To simplify matters, the auction sequence is fixed at first selling A, then B, then A etc. so that each synergy buyer will face exactly two auctions for an item of type A and two for an item of type B. However, the general result of this section is also true for a random auction sequence, since the basic effects remain the same.




Figure 5.10: Percentage increase in seller's profits in a larger experimental setting, with synergy buyers randomly entering the market.

As shown in Figure 5.10, the seller's profit now only has one maximum at 5, because initially each increase in exercise prices causes, with some probability, a synergy buyer to

participate more often. So each point is a threshold and the profit graph smooths out over those many local maxima, corresponding to a steady increase (on average) of the expected profit. This result shows why it can be rational for the seller to have the same exercise prices for all goods of the same type (e.g. the same K_A). In a market with random entry of synergy buyers, the seller does not know which buyers are participating in any particular auction. Her optimal policy is to set her exercise prices which maximize her overall expected profit (in this case, $K = 5$).

5.7 Discussion and further work

This chapter examined, from a decision-theoretic perspective, the use of priced options as a solution to the exposure problem in sequential auctions. We consider a model in which the seller is free to fix the exercise price for options on the goods she has to offer, and then sell these options in the open market, through a regular auction mechanism.

For this setting, we derived analytically, for a market with a synergy buyer and under some assumptions, the expressions that provide the bounds on the option prices between which both synergy buyers and sellers have an incentive to use an option contract over direct auctions. Next, we performed an experimental analyses of several settings, where either one or multiple synergy bidders are active simultaneously in the market. We show that, if the exercise price is chosen correctly, selling items through priced options rather than direct sale can increase the expected profits of both parties.

The overall conclusion of our study is that the proposed priced options mechanism can considerably reduce the exposure problem that synergy bidders face when taking part in sequential auctions. Furthermore, and most important, *both* parties in the market have an incentive to prefer and use such a mechanism. We show that in many realistic market scenarios, sellers can fix the exercise prices at a level that both provides sufficient incentive for buyers to take part in the auctions, as well as cover their risk of remaining with the items unsold.

We stress, however, that sequential auction allocation and bidding is a highly complex and still under-researched area, and our study provides just a first decision-theoretic analysis for the use of options to solve this problem. Basically, we provide the analysis and results for several fundamental cases, which can serve as a basis for future work on other settings. These include more complex market scenarios, as well as more sophisticated reasoning abilities on the part of participating synergy bidders and sellers. For example, in a large market, synergy bidders could be expected to use learning strategies to adapt to changing market conditions, as well as the presence of other synergy bidders who want similar item combinations. However, the sellers of the items could also use learning to choose better levels of the exercise prices K with which to sell the options for their goods. Other possible issues open to future research include: markets where bidders have asymmetric or imperfect information, more complex preferences over bundles and different attitudes to risk of the involved parties.

To conclude, sequential auction bidding with complementary valuations is a problem that appears in many real-life settings, although no dominant strategies exist and bidders face a severe exposure problem. The main intuition of this work is that a simple options mechanism, where sellers auction options for their goods (with a pre-set exercise price), instead of the goods themselves can go a long way in solving the exposure problem, and can be beneficial to both sides of such a market.

In practical terms, the potential impact of having a working solution to the exposure problem in sequential auctions is considerable. For example, in transportation logistics (see the study case in Appendix A of this thesis) many loads appear sequentially, over time, and a bidding agent has to acquire a combination of these to fill her transportation capacity (i.e. truck). In decentralised electricity markets, much of the available electricity supply (especially that generated by renewable sources, such as wind or solar energy) comes online with a certain probability. In allocating this intermittent, “green” electricity through a market-based method, options could be a promising solution to deal with the inherent uncertainty.

Other potential applications include retail electronic commerce (such as discussed in the follow-up work of Juda & Parkes [121]) or keyword markets in sponsored search (see case study in Chapter 7 of this thesis). We are planning to explore some of these potential applications in our future work.

Part III

Emergence of collaboration and social preferences in web communities and online markets

Chapter 6

Emergence of Consensus and Shared Vocabularies in Collaborative Tagging

6.1 Introduction

The previous chapters of this thesis presented experimental research relating to modeling preferences and in agent-mediated bilateral negotiation and auction settings. This chapter is the first of a series of two chapters that takes a different perspective on the issue of preference, by looking at how preferences form in large-scale web communities and online advertising markets. The methodology used differs in these two chapters: while Chapters 2 - 5 validated their conclusions mostly through computer simulations, this chapter and Chapter 7 use empirical analysis of large-scale web data, data produced by the actions of many (often millions) of actual web users.

The methods we found most useful to examine such large-scale, decentralized systems are those inspired by complex systems theory [11, 41, 166]. As we detail later in these two chapters, complex systems research aims to explore how order and structure can emerge in a system composed of many autonomous entities acting independently, without a central controller (such as a social web community or an online market). Examining the dynamics of economic phenomena using agent-based computational economics (ACE) methods [218] is a prominent example of such an approach.

This chapter focuses on one class of web systems that exhibits such phenomena and which has recently received a lot of interest in the world-wide web community: collaborative tagging. This work resulted after a collaboration, based on an extended stay at the Santa Fe Institute, Santa Fe, NM. While it is true that, for tagging systems, one cannot talk of

preferences in the economic, utility-based sense, we argue this work is closely related to the main topic of this thesis, modeling preferences and decisions in agent-mediated electronic markets. Actions such as clicking on a link or choosing a tag can be implicitly seen as expressing a preference, and our goal is to examine the dynamics of this process, as well as the information structures that emerge from it. In the next chapter, Chapter 7, we use complex systems theoretic methods to empirically explore some important properties of sponsored search advertising markets.

For both chapters, the most important aspect we follow is how a collective consensus (defined here in terms of tag, respectively keyword distributions) can emerge from the decentralized actions of a large number of agents (in this case, users of the system). Furthermore, both chapters make use of graphical and similarity-based techniques, related to - and initially inspired from - the ones developed in Chapter 3 of this thesis.

6.1.1 Tagging versus Taxonomies on the Web

The issue of how knowledge engineering on the Web should proceed with the greatest efficiency and efficacy is a central concern as the amount of information on the Web grows. A small but increasingly influential set of web applications, including the social bookmarking site del.icio.us, Flickr, Furl, Rojo, Connotea, Technorati, and Amazon allow users to *tag* objects with keywords to facilitate retrieval both for the acting user and for other users. Sets of categories derived based on the tags used to characterize some resource are commonly referred to as folksonomies. This approach to organizing online information is usually contrasted with taxonomies, including the approach some associate with the Semantic Web.

There are concrete benefits to the tagging approach. The flexibility of tagging systems is thought to be an asset; tagging is considered a categorization process, in contrast to a pre-optimized classification process such as expert-generated taxonomies. In defining this distinction, [110] believes that “categorization divides the world of experience into groups or categories whose members share some perceptible similarity within a given context. That this context may vary and with it the composition of the category is the very basis for both the flexibility and the power of cognitive categorization.” Classification, on the other hand “involves the orderly and systematic assignment of each entity to one and only one class within a system of mutually exclusive and non-overlapping classes; it mandates consistent application of these principles within the framework of a prescribed ordering of reality” [110].

Other authors argue that tagging enables users to order and share data more efficiently than using classification schemes; the free-association process involved in tagging is cognitively much more simple than are decisions about finding and matching existing categories [36]. Additionally, proponents of tagging systems show that users of tagging systems only need to agree on the general meaning of a tag in order to provide shared value instead of agreeing on a specific, detailed taxonomy [154].

However, a number of problems stem from organizing information through tagging sys-

tems including ambiguity in the meaning of tags and the use of synonyms which creates informational redundancy. Additionally, an important open question concerning the use of collaborative tagging to organize metadata is whether the system becomes *stable* over time. By *stable*, we mean that users have collectively developed some implicit consensus about which tags best describe a site, and these tags do not vary much over time. We will assume that these tags that best describe a resource will be those that used most often, and new users mostly reinforce already-present tags with similar frequencies. Since users of a tagging system are not acting under a centralized controlling vocabulary, one might imagine that no coherent categorization schemes would emerge at all from collaborative tagging. In this case, tagging systems, especially those with an open-ended number of non-expert users like del.icio.us, would be inherently unstable such that the tags used and their frequency of use would be in a constant state of flux. If this were the case, identifying coherent, stable structures of collective categorization produced by users with respect to a site would be difficult or impossible.

Given the debate over the utility of collaborative tagging systems compared to other methods of knowledge engineering on the Web, it is increasingly important to understand whether a coherent and socially navigable method of categorization can emerge from collaborative tagging systems. This paper will empirically examine a crucial aspect of this question: whether tag distributions stabilize over time and, if so, what type of distributions emerge. Since each tag for a given web resource (such as a web page) is repeated a number of times by different users, for any given tagged resource there is a distribution of tags and their associated frequencies. The collection of all tags and their frequencies ordered by rank frequency for a given resource is the *tag distribution* of that resource.

The hope among proponents of collaborative tagging systems is that stable tag distributions, and thus, possibly, stable categorization schemes, might arise from these systems. Again, by *stable* we do not mean that users stop tagging the resource, but instead that users collectively settle on a group of tags that describe the resource well and new users mostly reinforce already-present tags with the same frequency as they are represented in the existing distribution. Online tagging systems have a variety of features that are often associated with complex systems such as a large number of users and a lack of central coordination. These types of systems are known to produce a distribution known as a power law over time. A crucial feature of some power laws - and one that we also exploit in this work - is that they can be produced by scale-free networks. So regardless of how large the system grows, the shape of the distribution remains the same and thus *stable*. Researchers have observed, some casually and some more rigorously, that the distribution of tags applied to particular resources in tagging systems follows a power law distribution where there are a relatively small number of tags that are used with great frequency and a great number of tags that are used infrequently [154]. If this is the case, tag distributions may provide the stability necessary to draw out useful information structures.

This chapter empirically examines two important questions regarding the structure of tagging systems; first, whether tag distributions stabilize over time, and if so, what type of distribution emerges and second, whether the resulting structure of tags can be utilized to

construct categorizations that provide meaningful information. This work seeks to make a contribution both to the theoretical understanding of the nature of tagging systems and to applied problems of information extraction from tagging systems.

6.1.2 Overview of Related Work

Existing research on tagging has explored a wide variety of problems, ranging from fundamental to more practical concerns. In this section, we provide a broad overview of the types of problems that interest researchers and practitioners in this area. We then focus on the research most relevant to the work presented here, in order to underscore our contribution.

A number of papers (Halvey & Keane '07 [94], Hearst & Rosner '08 [97], Byron et. al. '07 [132]) examine which tag presentation techniques enable users to find information with greatest ease and speed. They often put a special emphasis on tag clouds, the most widely used presentation technique). Halvey and Keane [94] provide a systematic evaluation of the properties of tag interfaces which have the most effect on the accuracy and speed with which users find information. Using a set of 62 test subjects, they show that alphabetization, font size and position of the tags play an important role. They also conclude that users scan lists and clouds of tags, rather than reading them directly. Byron et. al. [132] perform a similar study, but focused on the field of biomedical information. They compare the results of user search based on the PubMed database with results from a search using tag clouds extracted from search summaries returned by PubMed. They conclude that a tag cloud interface is advantageous in presenting descriptive information, but it may be less effective in enabling users to discover relationships between concepts than full text summaries.

In more recent work, Hearst and Rosner '08 [97] extend the study of tag clouds by also examining the subjective reactions of the test users to different layouts. They also discuss the role that social signaling may play in motivating the use of tag clouds. Another paper concerned with visualization is Kaser and Lemire '07 [123], who study the performance of different visualization algorithms for the 2-dimensional tag cloud drawing problem. The algorithms proposed are evaluated based on criteria such as minimization of the screen area required and computational speed. Compared to our work, this direction of research on tag visualization is different in scope, since we are more concerned with macro-level properties of tagging systems (e.g. convergence, emergence of shared vocabularies) than with visualization and usability aspects. However, as future work, comparing visualization methods using tag correlation graphs (as discussed in Sect. 6.4 of this chapter) with existing approaches using tag clouds may prove insightful.

Boydell and Smyth '06 [28] propose an approach for building a community-based snippet index that reflects the expertise and revolving interests of a group of searchers. They show how such an index could be used to re-rank the results produced by an underlying search engine, such as to give a higher rank to results that have been frequently selected by members of the same community in the past. Boydell and Smyth '07 [29] build on the idea of using community knowledge, by proposing a social summarization technique which allows the generation of more community-focused and query-sensitive summaries than those

returned by standard search engines. While this line of work does not focus explicitly on tagging, it uses the same underlying principle, that of capturing the expertise of a community of like-minded searchers to improve search results.

Another line of research is concerned with understanding how individual users perform tagging, and assisting them in choosing more useful tags. Sood et. al. '07 [196] propose a method to assist the users of a tagging system by suggesting tags potentially useful to other users based on existing tag posts. Kelkar and Seligmann '07 [203] propose two measures of the value of tags, the intensity and the spread and discuss how these could be used in information retrieval. Berendt and Hanser '07 [15] examine the motivations of users of a tagging system and argue that users see tagging not as a way of adding metadata to resources, but simply as adding "more content". Zollers '07 [235] additionally attempts to characterize the motivations of users with data from Amazon.com and Last.fm. She identifies the expression of an opinion, performance (i.e. self-presentation), and online activism as three main motivations. Michlmayr and Cayzer '07 [155] present a method for creating user profiles from tagging data and then leveraging them for personalized information access. Using the results from a small-scale user study they show that tag co-occurrence information can more successfully learn personalized user profiles than can single tag frequencies.

Other research examines the use of tagging for specific contexts and applications. Hayes and Avesani '07 [95] provide a discussion of how tag clustering techniques could be used to retrieve information in blogs, while Bateman et. al. '07 [13] describe how using tagging in an e-learning system can supplement traditional metadata-gathering approaches. Dubinko et. al. [65] consider the problem of visualizing the evolution of tags within the Flickr community. They develop several methods and algorithms for dynamically presenting tags to users given a sliding time window. Rattenbury, Good & Naaman '07 [182] present a method for the automatic extraction of event and place semantics from Flickr tags. [52] develop a system for the automatic generation of personalized tags for webpages during browsing. The tags/keywords generated are based both on the textual content of the webpage being browsed and on the data residing on the surfer's desktop. All of these techniques would benefit from a method for determining whether a given set of tags has stabilized, such as the one proposed in this chapter, in order to present the most stable tags to the user or program. If tags were presented before they stabilized, the information presented to the user might be less valuable.

In a direction of work that bears directly on the larger question of this research, Mika '05 [156] addresses the problem of extracting taxonomic information from tagging systems in the form of Semantic Web ontologies. The chapter extends the traditional model of taxonomies by incorporating a social dimension, thus establishing an essential connection between tagging and the techniques developed in the Semantic Web arena. However, unlike this work, Mika does not study the stabilization of the tag distributions themselves. Ideally, one would want to know if a tag distribution was stable before attempting to extract any taxonomic information from it.

There are several lines of research which take a perspective closely related to our work. Shen and Wu '07 [202] are interested in the structure of a tagging network for del.icio.us data as we are in Section 6.4. Unlike in our examples, their graph is unweighted [202] and

does not reflect the information in the tag distribution. They examine the degree distribution (the distribution of the number of other nodes each node is connected to) and the clustering coefficient (based on a ratio of the total number of edges in a subgraph to the number of all possible edges) of this network and find that the network is indeed “scale-free” and has the features Watts and Strogatz [229] found to be characteristic of small world networks: small average path length and relatively high clustering coefficient. A considerable amount of work exploring the structural properties of natural language networks finds similar results (Cancho and Sole ’03 [40]).

Another paper that studies the meta-level properties of large folksonomy graphs (that appeared concurrently with the conference version of our paper) is Schmitz et. al. ’07 [199]. They provide an thorough, in-depth analysis of the statistical properties of folksonomy networks, studying, among others: their characteristic path length, their “cliquishness”, connectedness, as well as the associative behavior of nodes in such networks. However, unlike our work, they do not provide an insight into how such folksonomy networks could be used for actual information retrieval and visualization, or to build simple information structures such as shared vocabularies.

An early line of research that has attempted to formalize and quantify the underlying dynamics of a collaborative tagging systems is Golder and Huberman ’06 [86], which also make use of del.icio.us data. They show the majority of sites reach their peak popularity, the highest frequency of tagging in a given time period, within ten days of being saved on del.icio.us (67% in their data set), though some sites are “rediscovered” by users (about 17% in their data set), suggesting stability in most sites but some degree of “burstiness” in the dynamics that could lead to cyclical patterns of stability characteristic of chaotic systems. Importantly, Golder and Huberman find that the distribution of tags within a given site stabilizes over time, usually around one hundred tagging events. They do not, however, examine what type of distribution arises from a stabilized tagging process, nor do they present a method for determining the stability of the distribution which we see as central to understanding the possible utility of tagging systems.

In a very recent line of research, Heymann et. al. ’08 [101] provide a large-scale comparison between social bookmarking and traditional web search, also using del.icio.us data. They find that tags used on del.icio.us are, on the whole accurate, while the class of users that use this system is broad, i.e. not restricted to a small subset of users. They also observe, however, that a large proportion of the tags assigned to a webpage (or resource) already appear in the title, forward and backward links to that page. Therefore, while tags assigned to resources are accurate, their distributions may not be suitable to make a significant impact on search performance. This is somewhat in line with our findings: while tags converge relatively fast to stable, power law distributions (c.f. Sect. 2), the top of these distributions may contain common (or obvious) tags. A solution to this problem (also suggested in [101] may be a better mechanism for recommending tags. Conceivably, the local “vocabulary extraction” methods presented in Sect. 6.5 of this chapter (and adaptations thereof) could be used to this end.

One important result is represented by Cattuto et. al. ’07 [43], which discuss generative

models to produce power law distributions for tag correlations. They also take a complex systems perspective to tagging and propose a generic model for the behavior of taggers, in the form of a Yule-Simon process with memory. However, Cattuto et. al. do not provide an analysis of how tag frequencies per website actually converge in time to stable distributions. Dellschaft and Staab '08 [63] proposes a more-parametrized model that accounts for power law distributions in tag vocabulary growth and in tag distributions for websites. Overall, we see our work and that of Cattuto et. al. and Dellschaft and Staab as complementary in scope. While they provide a theoretical model of a process which could give rise to power law distributions in tagging, we propose using an information-theoretic technique in Section 6.3 to analyze the convergence of power law distributions in already-existing tagging systems. Furthermore, we demonstrate its utility in several applications, such as tag graphs, shared vocabularies, as shown in this chapter.

Yet another important recent direction of work is represented by Sen et. al. '06 [200]. They present a user-centric model of tagging that distinguishes between personal tendency and community influence in the behavior of individual taggers. Furthermore, they propose a method to select tags to be displayed to a user, such as to maximize tag utility, adoption and user satisfaction. By contrast to [200], we focus on the aggregate tag distributions per resource which, in a large tagging system are highly unlikely to be personal bookmarks, but rather reflect the opinion or consensus of the user community.

Finally, based on the empirical results presented in the WWW'07 conference version of this work [93], Mikroyannidis '07 [157] argues that Semantic Web and Social Web approaches are essentially compatible and can co-exist. While we agree with the arguments presented by Mikroyannidis [157], we should point out that convergence to stable tag distributions does not, by itself, imply that the converged distributions are directly usable for information retrieval. The process of constructing proper formal ontologies from folksonomies, while perhaps possible under certain conditions, is not a straightforward task.¹ First, one would not want to attempt to construct any taxonomy unless one had a reliable method for determining whether or not the tagging process had stabilized, such as the one we propose in Section 6.3. Also, we view the automated methods presented in Sect. 6.4 and 6.5 of this chapter, as fully automated first steps for any more formal taxonomy construction. While the shared tag vocabularies (c.f. Sect. 6.5 of this chapter) are not fully-fledged formal Semantic Web ontologies, they can also be useful structures for many information retrieval applications without any additional formalization.

6.1.3 The Tripartite Structure of Tagging

To begin, we review the conceptual model of generic collaborative tagging systems theorized by [152,156] in order to make predictions about collaborative tagging systems based on empirical data and based on generative features of the model.

There are three main types of entities that compose any tagging system:

¹This may require for instance, some decision support in guiding the user, or a more structured design of the interface used to input the tags.

- The users of the system (people who actually do the tagging)
- The tags themselves
- The resources being tagged (in this case, the websites)

Each of these can be seen as forming separate spaces consisting of sets of nodes, which are linked together by edges (see Fig. 6.1). The first space, the *user space*, consists of the set of all users of the tagging system, where each node is a user. The second space is the *tag space*, the set of all tags, where a tag corresponds to a term (“music”) or neologism (“toread”) in natural language. The third space is the *resource space*, the set of all resources, where usually each resource is a website denoted by a unique URI.² A tagging instance can be seen as the two edges that links a user to a tag and that tag to a given website or resource. Note that a tagging instance can associate a date with its tuple of user, tag(s), and resource.




Figure 6.1: Tripartite graph structure of a tagging system. An edge linking a user, a tag and a resource (website) represents one tagging instance

From Figure 6.1, we observe that tags provide the link between the users of the system and the resources or concepts they search for.

This analysis reveals a number of dimensions of tagging that are often under-emphasized. In particular, tagging is often a *methodology for information retrieval*, much like traditional search engines, but with a number of key differences. To simplify drastically, with a traditional search engine a user enters a number of tags and then an automatic algorithm labels the resources with some measure of relevance to the tags *pre-discovery*, displaying relevant resources to the user. In contrast, with collaborative tagging a user finds a resource and then adds one or more tags to the resource manually, with the system storing the resource and the tags *post-discovery*. When faced with a case of retrieval, an automatic algorithm does not

²A URI is a “Universal Resource Identifier” such as <http://www.example.com> that can return a webpage when accessed. Some tagging based systems such as Spurl (<http://www.spurl.net>) store the entire document, not the URI, but most systems such as del.icio.us store only the URI. The resource space, in this definition, represents whatever is being tagged, which may or may not be websites per se.

have to assign tags to the resource automatically, but can follow the tags used by the user. The difference between this and traditional searching algorithms is two-fold: collaborative tagging relies on human knowledge, as opposed to an algorithm, to directly connect terms to documents before a search begins, and thus relies on the collective intelligence of its human users to *pre-filter* the search results for relevance. When a search is complete and a resource of interest is found, collaborative tagging often requires the user to tag the resource in order to store the result in his or her personal collection. This causes a *feedback cycle*. These characteristics motivate many systems like del.icio.us and it is well-known that feedback cycles are one ingredient of complex systems [11], giving further indication that a power law in the tagging distribution might emerge.

6.1.4 Organization of the chapter

This chapter is organized as follows. In the first part of the chapter, we examine how to detect the emergence of stable “consensus” distributions of tags assigned to individual resources. In Section 6.2 we demonstrate a method for empirically examining whether tagging distributions follows a power law distribution. In Section 6.3 we show how this convergence to a power law distribution can be detected over time by using the Kullback-Leibler divergence. We further empirically analyze the trajectory of tagging distributions before they have stabilized, as well as the dynamics of the “long tail” of tag distributions. In the second part of the chapter, we examine the applications of these stable power law distributions. In Section 6.4 we demonstrate how the most frequent tags in a distribution can be used in inter-tag correlation graphs (or folksonomy graphs) to chart their relation to one another. Section 6.5 shows how these folksonomy graphs can be (automatically) partitioned, using community-based methods, in order to extract shared tag vocabularies. Finally, Section 6.6 provides an independent benchmark to compare our empirical results from collaborative tagging, by solving the same problems using a completely different data set: search engine query logs. The chapter concludes with a discussion of future work.

6.2 Detecting Power Laws in Tags

This section uses data from del.icio.us to empirically examine whether intuitions regarding tagging systems as complex systems exhibiting power law distributions hold.

6.2.1 Power Law Distributions: Definition

A *power law* is a relationship between two scalar quantities x and y of the form:

$$y = cx^{\alpha} \tag{6.1}$$

where α and c are constants characterizing the given power law. Eq. 6.1 can also be written as:

$$\log y = \alpha \log x + \log c \quad (6.2)$$

When written in this form, a fundamental property of power laws becomes apparent; when plotted in log-log space, power laws are straight lines. Therefore, the most simple and widely used method to check whether a distribution follows a power law and to deduce its parameters is to apply a logarithmic transformation, and then perform linear regression in the resulting log-log space. In this chapter we used a more powerful regression method to derive α that minimizes the bias in the value of the exponent (see [166] for the technical details).

The intuitive explanation of power law parameters in the domain of tagging is as follows: c represents the number of times the most common tag for that website is used, while α gives the power law decay parameter for the frequency of tags at subsequent positions. Thus, the number of times the tag in position p is used (where $p = 1..25$, since we considered the tags in the top 25 positions) can be approximated by a function of the form:

$$Frequency(p) = \frac{c}{p^{-\alpha}} \quad (6.3)$$

where $-\alpha > 0$ and $c = Frequency(p = 1)$ is the frequency of the tag in the first position in the tag distribution (thus, it is a constant that is specific for each site/resource).

6.2.2 Empirical Results for Power Law Regression for Individual Sites

For this analysis, we used two different data sets. The first data set contained a subset of 500 “Popular” sites from del.icio.us that were tagged at least 2000 times (i.e. where we would expect a “converged” power law distribution to appear). The second data set considers a subset of another 500 sites selected randomly from the “Recent” section of del.icio.us. Both sections are prominently displayed on the del.icio.us site, though “Recent” sites are those tagged within the short time period immediately prior to viewing by the user and “Popular” sites are those which are heavily tagged in general.³ While the exact algorithms used by del.icio.us to determine these categories are unknown, they are currently the best available approximations for random sampling of del.icio.us, both of heavily tagged sites and of a wider set of sites that may not be heavily tagged.

The mean number of users who tagged resources in the “Popular” data set was 2074.8 with a standard deviation of 92.9, while the mean number of users of the “Recent” data set was 286.1 with a standard deviation of 18.2. In all cases, the tags in the top 25 positions in the distributions have been considered and thus all of our claims refer to these tags. Since the tags are rank-ordered by frequency and the top 25 is the subset of tags that are actually

³All data used in the convergence analysis was collected in the week immediately prior to 19 Nov 2006.

available to del.icio.us users to examine for each site, we argue that using the top 25 tags is adequate for this examination.

Results are presented in Figure 6.2. In all cases, logarithm of base 2 was used in the log-log transformation.⁴

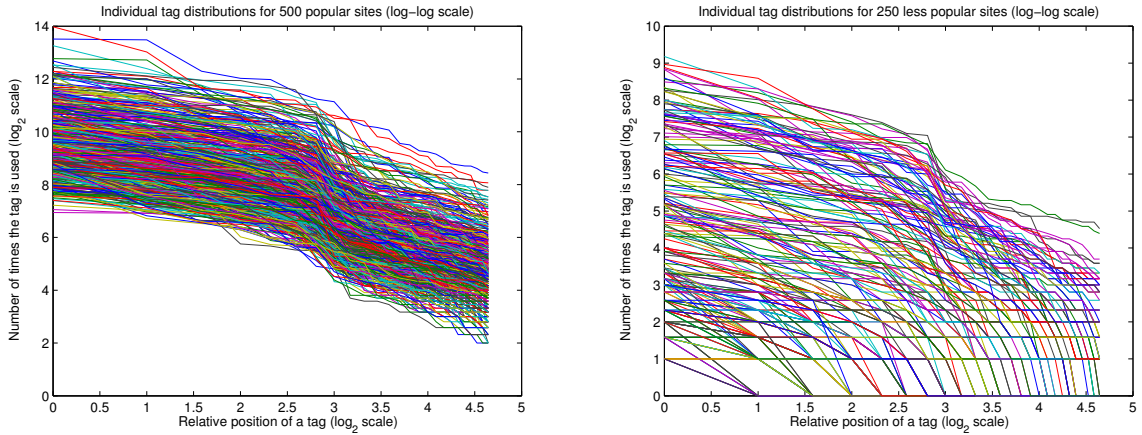


Figure 6.2: Frequency of tag usage relative to tag position. For each site, the 25 most frequently used tags were considered. The plot uses a double logarithmic (log-log) scale. The data is shown for a set of 500 randomly-selected, heavily tagged sites (left) and for a set of 500 randomly-selected, less-heavily tagged sites (right).

As shown by [166] and others, the main characteristic of a power law is its slope parameter α . On a log-log scale, the constant parameter c only gives the “vertical shift” of the distribution with respect to the y-axis. For each of the sites in the data set, the corresponding power law function was derived and the slopes of each (α parameters) were compared. The slopes indicate the fundamental characteristic of the power laws, as vertical shifts can and do vary significantly between different sites.

Our analysis shows that for the subset of heavily tagged sites, the slope parameters are very similar to one another, with an average of $\alpha = -1.22$ and a standard deviation ± 0.03 . Thus, it appears that the power law decay slope is relatively consistent across all sites. This is quite remarkable, given that these sites were chosen randomly with the only criteria being that they were heavily tagged. This pattern where the top tags are considerably more popular than the rest of the tags seems to indicate a fundamental effect of the way tags are distributed in individual websites which is independent of the content of individual websites. The specific content of the tags themselves can be very different from one website to the other and

⁴Note that the base of the logarithm does not actually appear in the power law equation (c.f. Eq. 6.1), but because we use empirical and thus possibly noisy data, this choice might introduce errors in the fitting of the regression phase. However, we did not find significant differences from changing the base of the logarithm to e or 10.

this obviously depends on the content of the tagged site.

For the set of less-heavily tagged sites, we found the slopes differed from each other to a much greater extent than with the heavily tagged data, with an average $\alpha = -5.06$ and standard deviation ± 6.10 . Clearly, the power law effect is much less pronounced for the less-heavily tagged sites as opposed to the heavily tagged sites, as the standard deviation reveals a much poorer fit of the regression line to the log-log plotted aggregate data. For sites with relatively few instances of tagging, the results reveal mostly noise.

6.2.3 Empirical Results for Power Law Regression Using Relative Frequencies

In the previous section, we applied power law regression techniques to individual sites, using the number of hits for a tag in a given position in the distribution. In this section, we examine the aggregate case where we no longer use the raw number of tags (because these are not directly comparable across sites), and instead use the relative frequencies of tags. The relative frequency is defined as the ratio between the number of times a tag in a particular position is used for a resource and the total number of times that resource is tagged⁵. Thus, relative frequencies for a given site always sum to one. These relative frequencies based on data from all 500 sites of the “Popular” data set were then averaged. Results are presented in Figure 6.3.

As before, a power law was derived in the log-log space using least-means squares (LMS) regression. This power law was found to have the slope $\alpha = -1.25$. The regression error, computed through the LMS method in the normal, not logarithmic space, was found to be 0.038. Note that the LMS regression error computation only makes sense when converted back in the normal space, since in the log-log space exponents are negative and, furthermore, deviations on the y-axis only denote actual error only after the \exp_2 function is applied. This corresponds to a LMS error rate in the power law regression of 3.8% over the total number of tags in the distribution, which is low enough to allow us to conclude that tag distributions do follow power laws.

We note, however, that there is a deviation from a perfect power law in the del.icio.us data in the sense that there is a change of slope after the top seven or eight positions in the distribution. This effect is also relatively consistent across the sites in the data set. This may be due to the cognitive constraints of the users themselves or an artifact of the way the del.icio.us interface is constructed, since that number of tags are offered to the users as a suggestion to guide their search process. Nevertheless, given that the LMS regression error is rather low, we argue the effect is not strong enough to change the overall conclusion that tag distributions follow power laws.

⁵To be more precise, the denominator is taken as the total number of times the resource is tagged with a tag from the top 25 positions, given available data.




Figure 6.3: Average relative frequency of tag usage, for the set of 500 “Popular” sites from above. On the y-axis, the logarithm of the relative frequency (probability) is given. (The plot uses a double logarithmic (log-log) scale, thus on the y-axis values are negative since relative frequencies are less than one.)

6.3 The Dynamics of Tag Distributions

In Section 6.2, we provide a method for detecting a power law distribution in the tags of a site or collection of sites. In this section, we study another aspect of the problem, namely how the shape of these distributions develops in time from the tagging actions of the individual users. First, we examine the how power law distributions form at the top (the first 25 positions) of tag distributions for each site. For this, we employ a method from information theory, namely the Kullback-Leibler divergence. Second, we study the dynamics of the entire tag distributions, including all tags used for a site, and we show that the relative weights of the top and tail of tag distributions converge to stable ratios in the data sets.

6.3.1 Kullback-Leibler Divergence: Definition

In probability and information theory, the Kullback-Leibler divergence (also known “relative entropy” or “information divergence”) represents a natural distance measure between two probability distributions P and Q (in our case, P and Q are two vectors representing discrete probability distributions). Formally, the Kullback-Leibler divergence between P and Q is defined as:

$$D_{KL}(P||Q) = \sum_x P(x) \log\left(\frac{P(x)}{Q(x)}\right) \quad (6.4)$$

The Kullback-Leibler distance is a non-negative, convex function, i.e. $D_{KL}(P, Q) \geq 0, \forall P, Q$ (note that $D_{KL}(P, Q) = 0$ iff. P and Q coincide). Also, unlike other distance measures it is not symmetric, i.e. in general $D_{KL}(P, Q) \neq D_{KL}(Q, P)$.

6.3.2 Application to Tag Dynamics

We use two complementary ways to detect whether a distribution has converged to a steady state using the Kullback-Leibler divergence:

- The first is to take the relative entropy between every two consecutive points in time of the distribution, where each point in time represents some change in the distribution. Again, in our data, tag distributions are based on the rank-ordered tag frequencies for the top 25 highest-ranked unique tags for any one website. Each point in time was a given month where the tag distribution had changed; months where there was no tagging change were not counted as time points. Using this methodology, a tag distribution that was “stable” would show the relative entropy converging to and remaining at zero over time. If the Kullback-Leibler divergence between two consecutive time points becomes zero (or close to zero), it suggests that the shape of the distribution has stopped evolving. This technique may be most useful when it is completely unknown whether or not the tagging of a particular site has stabilized at all.
- The second method involves taking the relative entropy of the tag distribution for each time step with respect to the final tag distribution, the distribution at the time the measurement was taken or the last observation in the data, for that site. This method is most useful for heavily tagged sites where it is already known or suspected that the final distribution has already converged to a power law.

The two methods are complementary; the first methodology would converge to zero if the two consecutive distributions are the same, and thus one could detect whether distributions converged if even temporarily. Cyclical patterns of stabilization and destabilization may be detected using this first method. The second method assumes that the final time point is the stable distribution so this method detects convergence only towards the final distribution. If both of these methods produce relative entropies that approach zero, then one can claim that the distributions have converged over time to a single distribution, the distribution at the final time point. Given our interest in distributions that have converged to power laws, we are actually examining the dynamics of convergence to a power law.

6.3.3 Empirical Results for Tag Dynamics

The analysis of the intermediate dynamics of tagging is considerably more involved than the analysis of final tag distributions. Because the length of the histories varies widely, there is no meaningful way to compute a cumulative measure across all sites as in Section 6.2, so

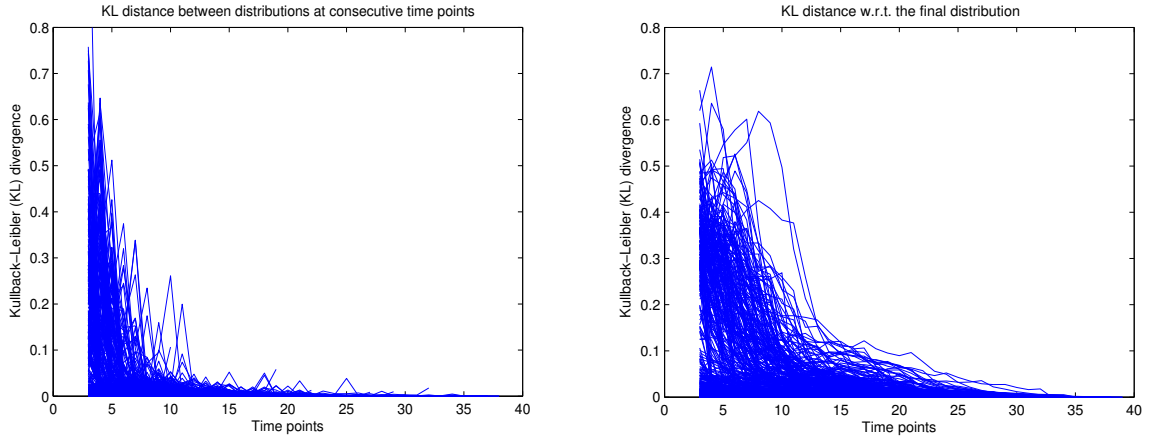


Figure 6.4: A (left). Kullback-Leibler divergence between tag frequency distributions at consecutive time steps for 500 “Popular” sites. B (right). Kullback-Leibler divergence of tag frequency distribution at each time step with respect to the final distribution.

our analysis has to consider each resource individually. In Figure 6.4 (A and B), we plot the results for the convergence of the 500 “Popular” sites, on the basis that their final distribution must have converged to a power law, that their complete tagging history was available from the first tagging instances, and that this history was of substantial length. In the data set considered, up to 35 time points are available for some sites (which roughly corresponds to three years of data, since one time point represents one month).

There is a clear effect in the dynamics of the above distributions.⁶ At the beginning of the process when the distributions contain only a few tags, there is a high degree of randomness, indicated by early data points. However, in most cases this converges relatively quickly to a very small value, and then in the final ten steps, to a Kullback-Leibler distance which is graphically indistinguishable from zero (with only a few outliers). If the Kullback-Leibler divergence between two consecutive time points (in Figure 6.4A) or between each step and the final one (Figure 6.4B) becomes zero or close to zero, it indicates that the shape of the distribution has stopped changing. The results here suggest that the power law may form relatively early on in the process for most sites and persist throughout. Even if the number of tags added by the users increases many-fold, the new tags reinforce the already-formed power law. Interestingly, there is a substantial amount of variation in the initial values of the Kullback-Leibler distance prior to the convergence. Future work might explore the factors underlying this variation and whether it is a function of the content of the sites or of the mechanism behind the tagging of the site. Additionally, convergence to zero occurs at approximately the same time period (often within a few months) for these sites.

⁶Note that in Figure 6.4, the first two time points were omitted because their distribution involved few tags and were thus very highly random.

The results of the Kullback-Leibler analysis provide a powerful tool for analyzing the dynamics of tagging distributions. This very well might be the result of the “scale-free” property of tagging networks, so that once the tagging of users have reached a certain threshold, regardless of how many tags are added, the distribution remains stable [202]. This method can be immensely useful in analyzing real-world tagging systems where the stability of the categorization scheme produced by the tagging needs to be confirmed.

6.3.4 Examining the dynamics of the entire tag distribution

In the previous sections, we focused on the distributions of the tags in the top 25 positions. However, heavily tagged or popular resources, such as those considered in our analysis, can be tagged several tens of thousands of times each, producing hundreds or even thousands of distinct tags. It is true that many of these distinct tags are simply personal bookmarks which have no meaning for the other users in the system. However, it is still crucial to understand their dynamics and the role they play in tagging, especially with respect to the top of the tag distribution. Some sources (e.g. Anderson [3]), have argued that the dynamics of long tails are a fundamental feature of Internet-scale systems. Here we were particularly interested in two questions. First, how does the number of times a site is tagged (including the long tail) evolve in time? Second, how does the relative importance of the head (top 25 tags) to the long tail change as tags are added to a resource?

Results for the same set of 500 “Popular” sites described above are shown in Figure 6.5. Note that the tag distributions were reconstructed through viewing the tagging history of the individual site as available through del.icio.us and collecting the growth of this tagging distribution over time, thus allowing us to record the growth of tags outside the 25 most popular.

As seen in Figure 6.5, the total number of times a site is tagged grows continuously at a rate that is specific to each site and this probably depends on its domain and particular context. Though the results are not shown here due to space constraints, a similar conclusion can be formulated for the number of distinct tags, given that the number of distinct tags varies considerably per site and does not seem to stabilize in time. However for virtually all of the sites in the data set considered, the proportion of times a tag from the top 25 positions is used relative to the total number of times that a resource is tagged did stabilize over time. So, while the total number of tags per resource grows continuously, the relative weight of the tags in the head of the tag distribution compared to the those in the long tail does stabilize to a constant ratio. This is an important effect and it represents a significant addition to our analysis of the stability analysis of the top 25 positions, since it shows the relative importance of the long tail with respect to the head of the distribution does eventually stabilize regardless of the growth of tags in the long tail.




Figure 6.5: A (left). Cumulative number of times a resource is tagged for each time point. B (right). Proportion of times a tag in the top 25 spots of the distribution has been used to tag a resource to the total number of times the resource has been tagged with any tag.

6.4 Constructing Tag Correlation Graphs

The previous section examines the type of frequency distributions that emerge from the collective tagging actions of individual users, as well as the dynamics of this process. This section examines the type of information structures that form from these actions, given the hypothesized importance of the information value of tags in understanding tagging systems. We look at one of the most simple information structures that can be derived through collaborative tagging: inter-tag correlation graphs (or, perhaps more simply, “folksonomy graphs”). We discuss the methodology used for obtaining such graphs and then illustrate our approach through an example domain study.

6.4.1 Methodology

The act of tagging resources by different users induces, at the tag level, a simple distance measure between any pair of tags. This distance measure captures a degree of co-occurrence which we interpret as a similarity metric, between the concepts represented by the two tags.

The collaborative filtering [185, 197] and natural language processing [151] literature proposes several distance or similarity measures that can be employed for such problems. The metric we found most useful for this problem is *cosine distance*.⁷

⁷This should not be interpreted as a conclusion on our part that cosine distance is always an optimal choice for this problem. This issue probably requires further research and even larger data sets.

Formally, let T_i, T_j represent two random tags. We denote by $N(T_i)$ and $N(T_j)$ respectively the number of times each of the tags was used individually to tag all resources, and by $N(T_i, T_j)$ the number of times two tags are used to tag the same resource. Then the similarity between any pair of tags i and j is defined as:

$$\text{similarity}(T_i, T_j) = \frac{N(T_i, T_j)}{\sqrt{N(T_i) * N(T_j)}} \quad (6.5)$$

In the rest of the chapter, we use the shorthand: sim_{ij} to denote $\text{similarity}(T_i, T_j)$.

From these similarities we can construct a tag-tag correlation graph or network, where the nodes represent the tags themselves weighed by their absolute frequencies, while the edges are weighed with the cosine distance measure. We build a visualization of this weighed tag-tag correlation, by using a “spring-embedder” or “spring relaxation” type of algorithm. We tested two such algorithms: Kawada-Kawai and Fruchterman-Reingold [12]; the two graphs included in this chapter are based on the latter. An analysis of the structural properties of such tag graphs may provide important insights into both how people tag and how structure emerges in collaborative tagging.

6.4.2 Constructing the tag correlation (folksonomy) graphs

In order to exemplify our approach, we collected the data and constructed visualizations for a restricted class of 50 tags, all related to the tag “complexity.” Our goal in this example was to examine which sciences the user community of del.icio.us sees as most related to “complexity” science, a problem which has traditionally elicited some discussion. The visualizations were made on Pajek [12]. The purpose of the visualization was to study whether the proposed method retrieves connection between a central tag “complexity” and related disciplines. We considered two cases:

- Only the dependencies between the tag “complexity” and all other tags in the subset are taken into account when building the graph (Fig. 6.6).
- The weights of all the 1175 possible edges between the 50 tags are considered (Fig. 6.7).

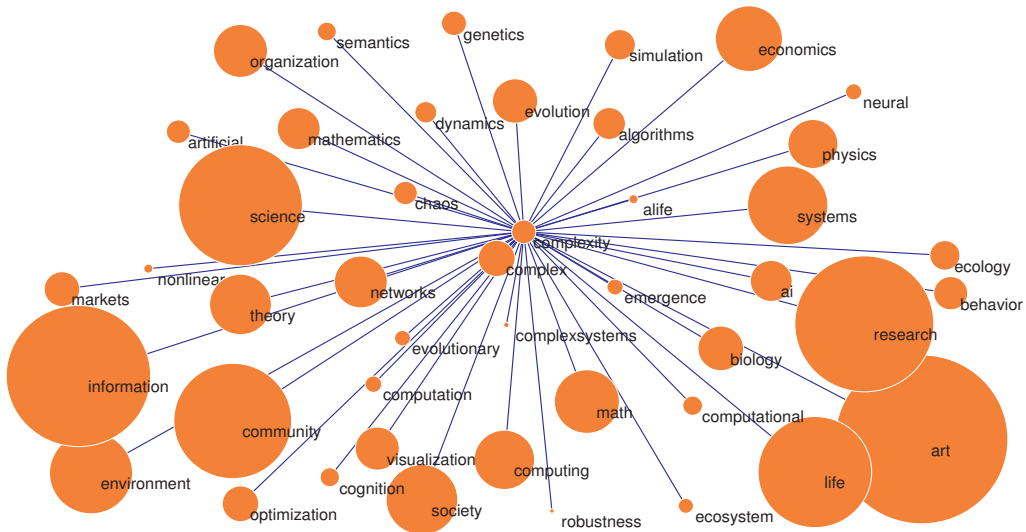


Figure 6.6: Folksonomy graph, considering only correlations corresponding to central tag “complexity”

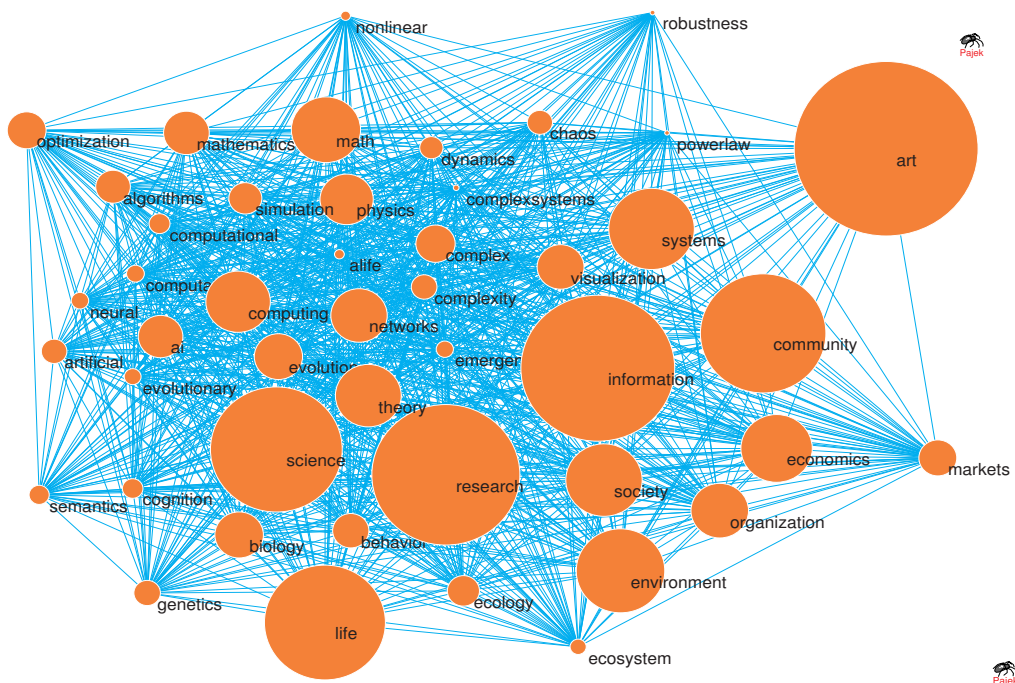


Figure 6.7: Folksonomy graph, considering all relevant inter-tag correlations

In both figures, the size of the nodes is proportional to the absolute frequencies of each tag, while the distances are, roughly speaking, inversely related to the distance measure as returned by the “spring-embedder” algorithm.⁸ We tested two energy measures for the “springs” attached to the edges in the visualization: Kamada-Kawai and Fruchterman-Reingold [12]. For lack of space, only the visualization returned by Kamada-Kawai is presented here, since we found it more faithful to the proportions in the data.

The results from the visualization algorithm match relatively well with the intuitions of an expert in this field. Some nodes are much larger than others which again shows that taggers prefer to use general, heavily used tags (e.g. the tag “art” was used 25 times more than “chaos”). Tags such as “chaos”, “alife”, “evolution” or “networks” which correspond to topics generally seen as close to complexity science are close to it. At the other end, the tag “art” is a large, distant node from “complexity.” This is not so much due to the absence of sites discussing aspects of complexity in art as there are quite a few of such sites, but instead due to the fact that they represent only a small proportion of the total sites tagged with “art,” leading to a large distance measure.

In Figure 6.7, the distances to “complexity” change significantly, due to the addition of the correlations to all other tags. However, one can observe several clusters emerging which match reasonably well with intuitions regarding the way these disciplines should be clustered. Thus, in the upper-left corner one can find tags such as “mathematics”, “algorithms”, “optimization”, “computation”, while immediately below are the disciplines related to AI (“neural” [networks], “evolutionary” [algorithms] and the like). The bottom left is occupied by tags with biology-related subjects, such as “biology”, “life”, “genetics”, “ecology” etc, while the right-hand side consists of tags with more “social” disciplines (“markets”, “economics”, “organization”, “society” etc.). Finally, some tags are both large and central, pertaining to all topics (“research”, “science”, “information”).

We also observed some tags that are non-standard English words, although we filtered most out as not relevant to this analysis. One example is “complexsystems” (spelled as one word), which was kept as such, although the tags “complex” and “system” taken individually are also present in the set. Perhaps unsurprisingly, the similarity computed between the tags “complexsystems” and “complex” is one of the strongest between any tag pair in this set. One implication of this finding is that tag distances could be used to find tags that have minor syntactic variance with more well-known tags, such as “complexsystems,” but which cannot simply be detected by morphological stemming.

⁸For two of the tags, namely “algorithms” and “networks,” morphological stemming was employed. So both absolute frequencies and co-dependencies were summed over the singular form tag, i.e. “network” and the plural “networks,” since both forms occur with relatively high frequency.

6.5 Identifying tag vocabularies in folksonomies using community detection algorithms

The previous sections analyzed the temporal dynamics of distribution convergence and stabilization in collaborative tagging as well as some information structures, like tag correlation (or folksonomy) graphs, that can be created from these tag distributions. In this Section, we look at how these folksonomy graphs could be used to solve an important problem in collaborative tagging: identifying shared tag vocabularies.

The problem considered in this section can be summarized as: given a heterogeneous set of tags (which can be represented as a folksonomy graph), how can we partition this set into subsets of related tags? In this chapter, we call this problem a “vocabulary identification” problem. It is important to note that we use the term “vocabulary” only in a restricted sense, i.e. as a collection of related terms, relevant to a specific domain. For instance, a list of tropical diseases is a “vocabulary”, a list of electronic components in a given electronic device is a vocabulary, and a list of specialized terms connected to a given scientific subfield would all be “vocabularies” in our definition.

We acknowledge that this is a restricted definition: in some applications, especially Semantic Web approaches, we would also like to know precisely how these terms are related. This type of structural information is difficult to extract only from tags, given the simple structure of folksonomies. Nevertheless, our approach could still prove useful in such applications: for example, one could construct the set of related terms as a first rough step and then a human expert (or, perhaps, another [semi]-automated method) could be used to add more detail to the extracted vocabulary set.

However, there are many settings in which the fully automated technique presented in this chapter could prove very useful. For example, drawing of tag clouds has received significant attention, but how to select the subset of related tags that will be presented in a cloud is an open problem. Another potential application is in selecting terms for sponsored search auctions. Some keywords (tags) bring a high value to advertisers, and knowing all the related keywords in a category that people can potentially use in search for can be very useful information for an advertiser. Conversely, the information regarding subsets of related tags could also be useful for the search engine in pricing searches using these tags.

Note that the complexity-related disciplines data set (already introduced in Sect. 4) is a useful tool to examine this question, since the initial set of tags are heterogeneous (complexity science is, by its very nature, an interdisciplinary field), but there are natural divisions into sub-fields, based on different criteria. This allows easier intuitive interpretation of the obtained results (besides the mathematical modularity criteria described below).

The technique we will use in our approach is based on the so-called “community detection” algorithms, developed in the context of complex systems and network analysis theory [164,165]. Such techniques have been well studied at a formal level and have been used to study large-scale networks in a variety of fields from social analysis (e.g. analysis of co-citation networks), analysis of biological nets (e.g. food chains) to gene interaction

networks. [165] provide an overview of existing applications of this theory, while [164] presents a formal analysis of the algorithm class used in this chapter. To the best of the authors' knowledge, however, this is the first work that studies the application of these techniques to tagging systems and folksonomies. In a somewhat related direction of work, [113] study the application of community detection techniques to aggregate bidder preferences in Ebay auctions.

6.5.1 Using community detection algorithms to partition tag graphs

In network analysis theory, a community is defined as a subset of nodes that are connected more strongly to each other than to the rest of the network. In this interpretation, a community is related to clusters in the network. If the network analyzed is a social network (i.e. vertexes represent people), then "community" has an intuitive interpretation. For example, in a social network where people who know each other are connected by edges, a group of friends are likely to be identified as a community, or people attending the same school may form a community. We should stress, however, that the network-theoretic notion of community is much broader, and is not exclusively applied to people. Some examples [113, 165] are networks of items on Ebay, physics publications on arXiv, or even food webs in biology. We will use a community detection algorithm to identify "vocabularies" within a folksonomy graph, identifying "communities" as "vocabularies."

Community detection: a formal discussion

Let the network considered be represented a graph $G = (V, E)$, when $|V| = n$ and $|E| = m$. The community detection problem can be formalized as a partitioning problem, subject to a constraint. The partitioning algorithm will result in a finite number of explicit partitions, based on clusters in the network, that will considered "communities."

Each $v \in V$ must be assigned to exactly one cluster C_1, C_2, \dots, C_{n_C} , where all clusters are disjoint, i.e. $\forall v \in V, v \in C_i, v \in C_j \Rightarrow i = j$.

Generally speaking, determining the optimal partition with respect to a given metric is intractable, as the number of possible ways to partition a graph G is very large. [164] shows there are more than 2^{n-1} ways to form a partition, thus the problem is at least exponential in n . Furthermore, in many real life applications (including tagging), the optimal number of disjoint clusters n_C is generally not known in advance.

In order to compare which partition is "optimal", the global metric used is *modularity*, henceforth denoted by Q . Intuitively, any edge that in a given partition has both ends in the same cluster contributes to increasing modularity, while any edge that "cuts across" clusters has a negative effect on modularity. Formally, let $e_{ij}, i, j = 1..n_C$ be the fraction of all edges in the graph that connect clusters i and j and let $a_i = \frac{1}{2} \sum_j e_{ij}$ be the fraction of the ends of edges in the graph that fall within cluster i (thus, we have $\sum_i a_i = \sum_{i,j} e_{ij} = 1$).

The modularity Q of a graph $|G|$ with respect to a partition C is defined as:

$$Q(G, C) = \sum_i (e_{i,i} - a_i^2) \quad (6.6)$$

Informally, so Q is defined as the fraction of edges in the network that fall within a partition, minus the expected value of the fraction of edges that would fall within the same partition if all edges would be assigned using a uniform, random distribution. These partitions are identified as communities by [165]. In tagging, each of these partitions is identified as a vocabulary.

As shown in [164], if $Q = 0$, then the chosen partition c shows the same modularity as a random division.⁹ A value of Q closer to 1 is an indicator of stronger community structure - in real networks, however, the highest reported value is $Q = 0.75$. In practice, [164] found (based on a wide range of empirical studies) that values of Q above around 0.3 indicate a strong community structure for the given network.

We will return shortly to define the algorithm by which this optimal partition can actually be computed, but first some additional steps are needed to link this formal definition to our tagging domain.

6.5.2 Edge filtering step

As shown in tag graph construction step above, for our data set the initial inter-tag graph contains $\binom{50}{2} = 1225$ pairwise similarities (edges), one for each potential tag pair. Most of these dependencies are, however, spurious as they represent just noise in the data, and our analysis benefits from using only the top fraction, corresponding to the strongest dependencies.

In this chapter, we make the choice to filter and use in further analysis only the top $m = k_d * n$ edges, corresponding to the strongest pairwise similarities. Here, k_d is a parameter that controls the density of the given graph (i.e. how many edges are there to be considered vs. the number of vertexes in the graph). In practice, we take values of $k_d = 1..10$, which for the tag graph we consider means a number of edges from 500 down to 50.

6.5.3 Normalized vs. non-normalized edge weights

The graph community identification literature [165] generally considers graphs consisting of discrete edges (for example, in a social network graph, people either know or do not know each other, edges do not usually encode a “degree” of friendship). In our graph,

⁹Note that Q can also take values smaller than 0, which would indicate that the chosen partition is worse than expected at random.

Algorithm 4 *GreedyQ Determination*: Given a graph $G = (V, E)$, $|V| = n$, $|E| = m$ returns partition $\langle C_1, \dots, C_{n_C} \rangle$

1. $C_i = \{v_i\}, \forall i = \overline{1, n}$
 2. $n_C = n$
 3. $\forall i, j, e_{ij}$ initialized as in Eq. 6.7
 4. repeat
 5. $\langle C_i, C_j \rangle = \text{argmax}_{C_i, C_j} (e_{ij} + e_{ji} - 2a_i a_j)$
 6. $\Delta Q = \max_{C_i, C_j} (e_{ij} + e_{ji} - 2a_i a_j)$
 7. $C_i = C_i \cup C_j, C_j = \emptyset$ //merge C_i and C_j
 8. $n_C = n_C - 1$
 9. until $\Delta Q \leq 0$
 10. $\text{max}Q = Q(C_1, \dots, C_{n_C})$
-

however, edges represent similarities between pairs of tags (c.f. Eq. 6.5). There are two ways to specify edge weights.

The non-normalized case assigns each edge that is retained in the graph, after filtering, a weight of 1. Edges filtered out are implicitly assigned a weight of zero.

The normalized case assigns each edge a weight proportional to the similarity between the tags corresponding to the ends. Formally, using the notations from Eq. 6.5 and 6.6 from above, we initialize the values e_{ij} as:

$$e_{ij} = \frac{1}{\sum_{ij} sim_{ij}} sim_{ij} \quad (6.7)$$

Where $\frac{1}{\sum_{ij} sim_{ij}}$ is a normalization factor, which assures that $\sum_{ij} e_{ij} = 1$.

6.5.4 The graph partitioning algorithm

Since we have established our framework, we can now formally define the graph partitioning algorithm. As already shown, the number of possible partitions for this problem is at least 2^{n-1} (e.g. for our 50 tag setting $2^{50} > 10^{15}$). Therefore, to explore all these partitions exhaustively would be clearly unfeasible. The algorithm we use to determine the optimal partition (Alg. 4) is based on [164], and it falls into the category of “greedy” clustering heuristics.

Informally described, the algorithm runs as follows. Initially, each of the vertexes (in our case, the tags) are assigned to their own individual cluster. Then, at each iteration of the algorithm, two clusters are selected which, if merged, lead to the highest increase in the modularity Q of the partition. As can be seen from lines 5-6 of Alg. 1, because exactly

Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5	Cluster 6	Cluster 7
computation optimization visualization physics mathematics math computational algorithms information computing theory	markets economics society community organization ecology ecosystem environment	semantics cognition neural ai alife artificial life behavior simulation research	powerlaw nonlinear complexsystems dynamics chaos emergence networks systems complex complexity	genetics biology evolution evolutionary science	robustness	art
Tags that increase modularity the most, if eliminated: theory, science, research, simulation, networks.						

Figure 6.8: Optimal partition in tag clusters (i.e. “communities”) of the folksonomy graph, when the top 200 edges are considered. This partition has a $Q=0.34$. After eliminating the 5 tags mentioned at the bottom, Q can increase to 0.43.




Figure 6.9: Modularity (Q-factor) and number of partitions obtained from applying community detection algorithms to the scientific disciplines data set

two clusters are merged at each step, it is easy to compute this increase in Q as: $\Delta Q = (e_{ij} + e_{ji} - 2a_i a_j)$ or $\Delta Q = 2 * (e_{ij} - a_i a_j)$ (the value of e_{ij} being symmetric). The algorithm stops when no further increase in Q is possible by further merging.

Note that it is possible to specify another stopping criteria in Alg. 1, line 9, e.g. it is possible to ask the algorithm to return a minimum number of clusters (subsets), by letting the algorithm run until n_C reaches this minimum value.

6.5.5 Graph partitioning: experimental results

The experimental results from applying Alg. 1 to our data set are shown in Fig. 6.9. In Fig. 6.8 we present a detailed “snapshot” of the partition obtained for one of the experimental configurations. There are several interesting features of the results.

First, it becomes clear that using normalized edge weights produces partitions with higher modularity than assigning all the top edges the same weight of 1. This was intuitively hypothesized by us, since edge weights represent additional information we can use, but it was confirmed experimentally. Second, we are clearly able to identify partitions with a modularity higher than around 0.3, which exhibit a strong community structure according to [165]. Yet perhaps the most noteworthy feature of the partitions is the rapid increase both in the modularity factor Q and in the number of partitions, as the number of edges filtered decreases (from left to right, in our figure).

The filtering decision represents, in fact, a trade-off. Having too many edges in the graph may stop us from finding a partition with a reasonable modularity, due to the high volume of “noise” represented by weaker edges. However, keeping only a small proportion of the strongest edges (e.g. 100 or 50 for a 50-tag graph, in our example), may also have disadvantages, since we risk throwing away useful information. While a high modularity partition can be obtained this way, the graph may become too “fragmented”: arguably, dividing 50 tags into 10 or 15 vocabularies may not be a very useful.

Note that it is difficult to establish a general rule for what a “good” or universally “correct” partition should be in this setting. For example, even the trivial partition that assigns each tag to its own individual cluster cannot be rejected as “wrong” but such a trivial partition would not be considered a useful result for most purposes. In this chapter we generally report the partitions found to have the highest modularity for the setting. However, for many applications, having a partition with a certain number of clusters, or some average cluster size, may be more desirable. The clustering algorithm proposed here (Alg. 1) can be easily modified to account for such desiderata, by changing the stop criteria in line 9.

Fig. 6.8 shows the solution with the highest modularity Q for a graph with 200 edges, in which 7 clusters are identified. This partition assigns tags related to mathematics and computer science to Cluster 1, tags related to social science and phenomena to Cluster 2, complexity-related topics to Cluster 4 etc., while “art” is assigned to its own individual cluster. This matches quite well our intuition, and its modularity $Q = 0.34$ is above (albeit close) to the theoretical relevance threshold of 0.3. In Section 6.6 we will compare this partition (as well as the entire tag graphs constructed in Section 6.4) against an independent benchmark that addresses the same problem, but based on a completely different data set: search engine query logs. However, first we briefly present a method that can further improve the modularity of the retrieved tag graphs.

Algorithm 5 *GreedyQ Elimination*: Given a partition C_1, \dots, C_{n_C} of graph $G = (V, E)$ removes all vertexes $v_i \in V$ that increase Q

1. repeat
 2. $v_i = \operatorname{argmax}_{v_i} [Q(\dots, C_k \setminus \{v_i\}, \dots) - Q(\dots, C_k, \dots)]$
 3. $\Delta Q = \max_{v_i} [Q(\dots, C_k \setminus \{v_i\}, \dots) - Q(\dots, C_k, \dots)]$
 where $v_i \in C_k$ // C_k is the partition of vertex i
 4. until $\Delta Q \leq 0$
-

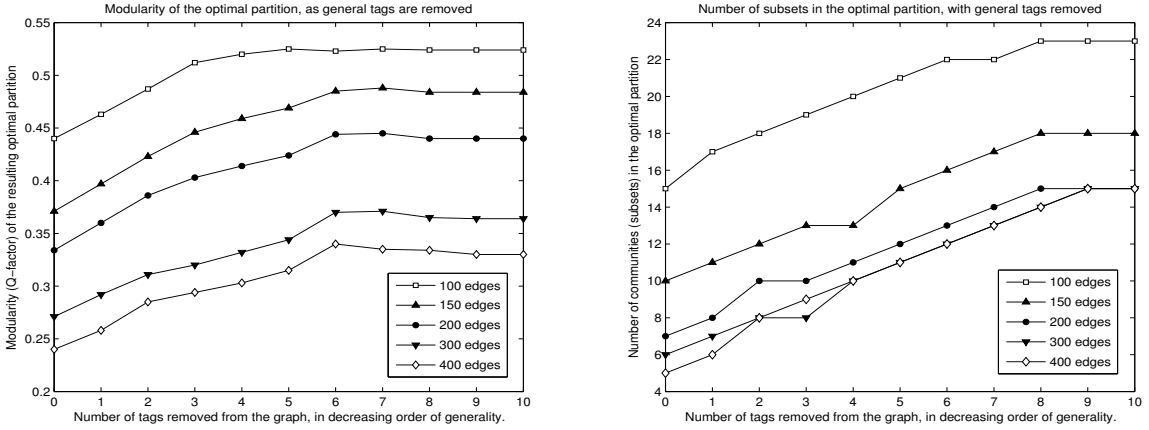


Figure 6.10: Modularity (Q-factors) and number of partitions obtained after gradually eliminating tags from the data set, such as to increase the modularity. At each step, the tag that produced the highest increase in modularity between the initial and resulting partition was selected. In these results, all edge weights are normalized.

6.5.6 Eliminating tags from resulting partitions to improve modularity

The analysis in the previous section shows that community detection algorithms were able to produce useful partitions, with above-relevance modularity. Still, there are a few general-meaning tags that would fit well into any of the subsets resulting after the partition. These tags generally reduce the Q modularity measure significantly, since they increase the inter-cluster edges. Therefore, we hypothesized that the modularity of the resulting partitions could be greatly improved by removing just a few tags from the set under consideration. In order to test this hypothesis, we tested another greedy tag elimination algorithm, formally defined as Alg. 2. Result graphs are shown in Fig 6.10, while in Fig. 6.8 we show the top 5 tags that, if eliminated, would increase modularity Q from 0.34 to 0.43.

As seen in Fig. 5, for this data set only 5-6 tags need to be eliminated as eliminating more does not lead to a further increases in Q . In the example in Fig. 6.8, we see which

these are, in order of elimination: theory, science, research, simulation, networks. In fact, these tags, that are marked for elimination automatically by Alg. 2, are exactly those that are the most general in meaning and would fit well into any of the subsets.

Regarding scalability, it is relatively straightforward to show that both Alg. 1 and 2 have linear running time the number of vertexes n , i.e. in this case, number of tags considered in the initial set. In the case of Alg. 1, exactly two clusters of tags are merged at each step, so one cluster increases in size by a minimum of one, until the algorithm terminates. In case of Alg. 2, one tag is eliminated per step, until termination. In practice, this scalability property means they are easily applicable to analyze much larger folksonomy systems.

To our knowledge, this is the first line of research to investigate the applicability of this type of algorithms to tagging, and we can conclude that results are very encouraging. We leave some aspects open to further work. For instance, in the current approach, similarity distances between pairs of tags are computed using all the tagging instances in the data set. In some applications, it might be useful to first partition the set of users that do the tagging, and then consider only the tags assigned by a certain class of users. For example, for tags related to a given scientific field, expert taggers may come up with a different vocabulary partition than novice users. This may require a two-fold application of this algorithm: first to partition and select the set of users, and then the set of tags based on the most promising category of users.

While these applications of tagging distributions have shown promise, one question that can be reasonably asked is how well these applications of tagging compare to some benchmark that does not use tagging distributions. In the next section we will compare the results obtained here from collaborative tagging data against a benchmark case, which uses “classic” search engine query data.

6.6 Comparison benchmark: automatic construction of key-word vocabularies from search engine query data

The previous sections of this chapter provide a compelling argument that show that a stable categorization scheme can arise from collaborative tagging, and these stable tagging distributions can produce vocabularies that can be harnessed in a wide range of applications. However, in order to truly establish the case for tagging, we need a benchmark to compare the results extracted from collaborative tagging data to results that can be obtained by means of other web search methods.

The obvious candidate for finding such a comparison benchmark is to use of large-scale query data produced by a search engine. The idea of approximating semantics by using search engine data has, in fact, been proposed before, and is usually found in existing literature under the name of “Google distance.” [53] were the first to introduce the concept of “Google distance” from an information-theoretic standpoint, while other researchers [85] have recently proposed using it for tasks such as approximate ontology matching. It is fair




Figure 6.11 Correlation graph from Microsoft queries, showing only correlations to the term “complexity”.

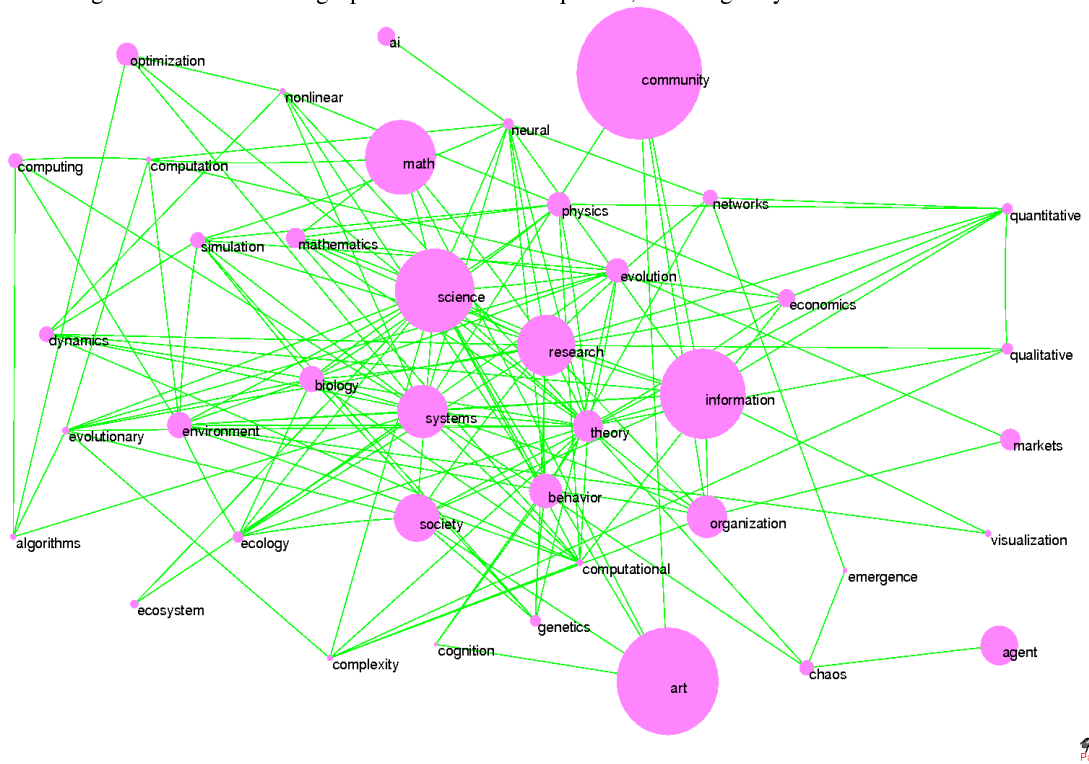


Figure 6.12: Correlation graph obtained from Microsoft query logs, considering all relevant search terms.

to assume (although we have no way of knowing this with certainty), that current search engines and related applications, such as Google Sets [106], also use text or query log mining techniques (as opposed to collaborative tagging) to solve similar problems.

There are two ways of comparing terms (in this case, keywords) using a search engine.

One method would be to compare the number of resources that are retrieved using each of the keywords and their combinations. Another method is to use the query log data itself, where the co-occurrence of the terms in the same queries vs. their individual frequency is the indicator of semantic distance. We employ this latter method as it is more amendable to comparison with our work on tagging. In the latter method, the query terms are comparable to tags, where instead of basing our folksonomy graphs and vocabulary extraction on tags, we used query terms. In general, query log data is considered proprietary and much more difficult to obtain than tagging data. We were fortunate to have access to a large-scale data set of query log data, from two separate proposals awarded through Microsoft’s “Beyond Search” awards.¹⁰ In the following we describe our methodology and empirical results.

6.6.1 Data set and methodology employed

The data set we used consists of 101,000,000 organic search queries, produced from Microsoft search engine Live.com, during a 3-month interval in 2006. Based on this set of queries, we computed the bilateral correlation between all pairs from the set of of complexity related terms considered in Sect. 6.4 and 6.5 above. The set of terms are, however, no longer treated as tags, but as search keywords.¹¹ The correlation between any two keywords T_i and T_j is computed using the cosine distance formula in Equation 6.5 from Section 6.4 above. However, here $N(T_i, T_j)$ represents the number of queries in which the keywords T_i and T_j appear in together, while $N(T_i)$ and $N(T_j)$ are the numbers of queries in which T_i , respectively T_j appear in total (irrespective of other terms in the query), from the 100 million queries in the data set.

The rest of the analysis mirrors closely the steps described in Sections 6.4 and 6.5, but optimizing the learning parameters which best fit this data set, in order to give both methods a fair chance in the comparison. More specifically, the Pajek visualization of the keyword graphs in Figs. 6.11 and 6.12 were also built by using a spring-embedder algorithm based on the Kamada-Kawai distance, while Fig. 6.13 shows the keyword vocabulary partition that maximizes the modularity coefficient Q in the new setting, considering the top 200 edges. For clarity, the graph pictures are depicted in a different color scheme, to clearly show they result from entirely different data sets: Figures 6.6 and 6.7 from del.icio.us collaborative tagging data, and Figures 6.11 and 6.12 from Microsoft’s Live.com query logs.

6.6.2 Discussion of the results from the query log data and comparison

When comparing the graphs in Figures 6.6 and 6.11 (i.e. the ones which only depict the relations to the central term “complexity”) an important difference can be observed. While the

¹⁰The authors wish to thank Microsoft Research for their kind support in providing this data.

¹¹We acknowledge this method has some drawbacks, as a few terms in the complexity-related set, such as “powerlaw” and “complexsystems” (spelled as one word) or “alife” (for “artificial life”) are natural to use as tags, but not very natural as search keywords. However, since there are only 3 such non-word tags, they do not significantly affect our analysis.

Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5	Cluster 6
complexity evolution evolutionary chaos cognition biology theory behavior markets genetics agent	systems visualization organization information community	networks ai emergence neural	algorithms ecology math computing optimization computation environment	mathematics physics economics art science simulation dynamics nonlinear computational ecosystem	research quantitative qualitative society
Terms left unclassified (i.e. one word clusters): complex, complexsystems, robustness, multi-agent, life, artificial, semantics, powerlaw, alife.					

Figure 6.13: Optimal partition into clusters, obtained from the Microsoft query data, when the top 200 edges are considered. The resulting partition has a $Q=0.536$. However, 9 terms were assigned to their own cluster, thus basically left unclassified.

graph in Fig. 6.6, based on collaborative tagging data, shows 48 terms related to complexity, the one is Fig. 6.11, based on query log data, shows just 6. The basic reason is that no relationship between the term “complexity” and the other 40+ terms can be inferred from the query log data. These relationships either do not appear in the query logs or are statistically too weak (only based on a few instances).

It is important to emphasize here that this result is not an artifact of the cosine similarity measure we use. Even if we use another, more complex distance measure between keywords, such as some suggested in the previous literature [53], we get very similar results. The fundamental reason for the sparseness of the resulting graph is the query log data itself does not contain enough relevant information about complexity-related disciplines. For example, among the 101,000,000 queries, the term complexity appears exactly 138 times, a term such as “networks” 1074 times. Important terms such as “cognition” or “semantics” are even less common, featuring only 47 and 26 times respectively among more than 100 million queries. Therefore, it is fair to conclude that the query log data, while very large in size, is quite poor in useful information about the complexity-related sciences domain. As a caveat, we do note that more common terms, such as “community” (78,862 times), “information” (36,520 times), “art” (over 52,000), or even “agent” (about 7,000) do appear more frequently, but these words have a more general language usage and are not restricted to the scientific domain. Therefore, these higher frequencies do not actually prove very useful for identifying the relationship of these terms to complexity science, which was our initial target question.

Turning our attention to the second graph in Fig. 6.12 and the partition in Fig. 6.13, we can see that query logs can also produce good results in comparison with tagging, although they are somewhat different from the ones obtained from tagging. For example, if we com-

pare the partitions obtained in Fig. 6.8 (resulting from tagging data) and the one in Fig. 6.13 (from query log data), we see that tagging produces a more precise partition of the disciplines into scientific sub-fields. For instance, it is clear from Fig. 6.8 that cluster 1 corresponds to mathematics, optimization and computation, cluster 2 to markets and economics, cluster 5 to biology and genetics, cluster 4 to disciplines very related to complexity science and so forth. The partition obtained from query log data in Fig. 6.13, while is still very reasonable, reflects perhaps how a general user would classify the disciplines, i.e. organization is related to both information, systems and community (cluster 2), research is either qualitative or quantitative (cluster 6), and the like. There are also some counter-intuitive associations, such as putting biology and markets in the same cluster (number 1). Note that the clustering (or modularity) coefficient Q is higher in Fig. 6.13 than 6.8, but this is only because there are less inter-connections between terms in general in the query log data, thus there are less edges to “cut” in the clustering algorithm.

To conclude, while both methods produce reasonable results, collaborative tagging does better, at least for this domain. Tagging data appears to be more rich in information about interconnections between the terms that can be exploited by the filtering algorithms proposed in this chapter. This can probably be explained by the fact the *del.icio.us* users have more expertise and interest in complexity-related topics than general web searchers. Furthermore, they are probably more careful in selecting resources to tag and in selecting labels for them that would be useful to other users as well (general web searchers are known to be “lazy” in typing queries. As a caveat, we note that this target domain (i.e. complexity-related disciplines) is scientific and very specialized. If the target would be more general (for example, if we selected a set of terms related to pop-culture), the comparison might lead to different results.

In future work, it may be interesting to study the formation of such vocabularies considering only the opinion (expressed in terms of bookmarks or queries) of a sub-community of users, such as the community of expert users employed in a particular field. While this should be theoretically possible for both approaches, in practice, it may be easier to trace identities of users with collaborative tagging, not least due to privacy concerns. People who sign up to use a collaborative tagging system are implicitly more willing to share their expertise with a group of users. By contrast, web search is a private activity, where tracing users’ expertise level or identity during search may be undesirable.¹²

6.7 Conclusions and Future Work

This work has explored the important question of whether a coherent, stable way of characterizing information can emerge from collaborative tagging systems and has presented several novel methods for analyzing data from such systems.

First, we have shown that tagging distributions of heavily tagged resources tend to stabilize into power law distributions and present a method for detecting power law distributions

¹²Although this probably happens, to some degree, in current practice.

in tagging data. We see the emergence of stable power law distributions as an aspect of what may be seen as collective consensus around some shared preferences regarding the categorization of information, consensus driven by tagging behaviors. We have additionally presented a method for examining the dynamics and convergence of stable tag distributions over time by the use of Kullback-Leibler divergence measures between distributions at different time steps. Also included is an empirical study of the importance of the “long tail” of the tag distributions in the convergence process.

In the second part of the chapter, we propose a method for constructing and visualizing correlation graphs from tags, and showed how they can lend important insights into how a community of users sees the relations between a set of terms. We also use a method from network theory for partitioning tag correlation graphs that can be used to identify vocabularies shared by a community of users. Finally, we show that vocabularies that from collaborative tagging data can be significantly richer, at least for some domains, than the ones that can be extracted from general search engine query logs. While these methods were empirically tested using del.icio.us data, the proposed methods are general enough to be applicable to other tagging systems.

This work suggests a number of exciting problems, both theoretical and applied, that merit further research. These include examining whether aspects of tagging distributions and dynamics are subject to the influence of particular features of tagging sites, to human cognitive limits, or some mixture of the two. A thorough examination of this aspect would represent a significant contribution to work in this area and would be important to many practical tagging applications.

Another important direction of work would be examining the effects of using specialized sub-communities of users in the study of convergence of tag distributions and resulting information structures, rather than the entire user population as in this chapter. As shown by [101], del.icio.us is not dominated by a small number of core users, but other tagging sites may be. We know relatively little about how user concentration might influence the types of information structures that can be derived from tags. Furthermore, the shared vocabulary used by a specialized sub-community of users may differ considerably to that of a larger user base.

Based on these results, it seems quite plausible that folksonomies can be fruitfully utilized for a wide category of applications related to organization of information on the web. Insights gained by taking collaborative tagging systems seriously as an empirical object of study could result in insight into the complexity of the one of the world’s most complex systems, the World Wide Web.

6.8 Acknowledgments

The initial stage of this work was performed during the authors’ stay at the Santa Fe Complex Systems Institute, Santa Fe, NM, USA. The authors thank SFI for its support.

Chapter 7

The Complex Dynamics of Sponsored Search Markets: An Empirical Study

7.1 Introduction

Sponsored search, the payment by advertisers for clicks on text-only ads displayed alongside search engine results, has become an important part of the Web. It promises to revolutionize advertising world and represents the main source of revenue for large search engines, such as Google, Yahoo! and Microsoft. Nevertheless, issues that arise from sponsored search also represent exciting research opportunities, in fields as diverse as economics, artificial intelligence, computer science and sociology.

For example, the field of multi-agent systems, researchers have been working for some time on topics such as designing automated auction bidding strategies in uncertain and competitive environments (Chapters 4 and 5 are two examples of this type of research, among many others, e.g. [19,230]). Another emergent field which studied such topic is agent-based computational economics (ACE), where significant research effort has focused on the dynamics of electronic markets through agent-based simulations. One particular topic of research for the ACE community is how order and macro-level market structure can emerge from the micro-level actions of individual users. However, most existing work has been based on simulations, as there are few sources of large-scale, empirical data from real-world automated markets. In this context, empirical data made available from sponsored search provides an excellent opportunity to test the assumptions made in such models in a real market.

In this paper, which is based on a large-scale Microsoft sponsored search dataset, we

provide such a detailed empirical analysis. To do this, we make use of several techniques derived from computational economics, and especially complex systems theory. Complex systems analysis (which we briefly review below) has been shown to be an excellent tool for analyzing large social, technological and economic systems, including web systems [41, 93, 166].

7.1.1 The data set

The study provided in this paper is based on a large dataset of sponsored search queries, obtained from the website Live.com¹. The search data provided consists of two distinct data sets: a set of sponsored search dataset (URLs returned are allocated to advertisers, through an auction mechanism) and an organic search dataset (standard, unbiased web search). The sponsored search data consists of 101,171,081 distinct impressions (i.e. single displays of advertiser links, corresponding to one web query), which in total received 7,822,292 clicks. This sponsored dataset was collected for a roughly 3-month period in the autumn of 2007. The organic search data set consists of 12,251,068 queries, and was collected in a different 3-month interval in 2006 (therefore the two data sets are chronologically disjoint).

It is important to stress that in the results reported in this paper are based mostly on the sponsored search data set². Furthermore, the sponsored search data we had available only provides partial information, in order to protect the privacy of Microsoft Live.com customers and business partners. For example, we have no information about financial issues, such as the prices of different keywords, how much different advertisers bid for these keywords, the budgets they allocate etc. Furthermore, while the database provides an anonymized identifier for each user performing a query, this does not allow us to trace individual users for any length of time.

Nevertheless, one can extract a great deal of useful information from the data. For example, the identities of the advertisers, for which keyword combinations their ads were shown (i.e. the impressions), for which of these combinations they received a click, the position their sponsored link was in when clicked etc. Insights gained from analyzing this information forms the main topic of this paper.

7.2 Complex systems analysis applied to the web and economics

Complex systems represents an emerging research discipline, at the intersection of diverse fields such as AI, sociology, economics and biology. The main focus of the study in the field of complex systems how macro-level dynamics may emerge from individual actions at

¹This data was kindly provided to us by Microsoft research through “Beyond Search” award

²The only exception is a plot on the distribution number of clicks vs. display rank in Sect. 7.3, included for comparison reasons.




Figure 7.1: Distribution of clicks received by a URL (link), relative to its position on the display, for sponsored and organic search. A(left-side, sponsored search dataset): There are up to 8 sponsored advertiser links displayed: 3 on the top of the page, and 5 in a side bar. B(right, organic search data): There are usually 10 positions displayed per page, with multiple result pages appearing as plateaus.

the individual level by agents participating in a system (such as an electronic market). For web phenomena, complex systems techniques have been successfully used before to study phenomena such as collaborative tagging [93], viral marketing [141] or the formation of online social groups [10].

One of the phenomena that are indicative to such complex dynamics is the emergence of scale-free distributions, such as power laws. The emergence of power laws in such a system usually indicates that some sort of complex feedback phenomena (e.g. such as a preferential attachment phenomena) is at work. This is usually one of the criteria used for describing the system as “complex” [11, 41]. Research in disciplines such as econophysics and computational economics discusses how such power laws can emerge in large-scale economic systems (see [41, 166] for a detailed discussion).

7.2.1 Power laws: definition

As a reminder from the previous Chapter, a power law is defined as a relationship between two scalar quantities x and y of the form:

$$y = cx^\alpha \quad \log y = \alpha \log x + \log c \quad (7.1)$$

where α and c are constants characterizing the given power law.

Power laws have the important property that when plotted in log-log space, power laws appear as straight lines. As shown by Newman [166] and others, the main parameter that characterizes a power law is its slope parameter α . (On a log-log scale, the constant parameter c only gives the “vertical shift” of the distribution with respect to the y-axis.). Vertical shift can vary significantly between different phenomena measured (in this case, click distributions), which otherwise follow the same dynamics. Furthermore, since the logarithm is applied to both sides of the equation, the size of the parameter α does not depend on the basis chosen for the of the logarithm (although the shifting constant c is affected). In the log-log plots shown in this chapter, we have chosen the basis of the logarithm to be 2, since we found graphs with this low basis the more graphically intuitive. But, in principle, the same conclusions should hold if we choose the logarithm basis to be, e.g. e or 10.

7.3 Influence of display rank on clicking behaviour

The first issue that we studied (for both sponsored and organic search data) is how the position that a URL link is displayed in influences its chances of receiving a click. Note that this particular issue has received much attention in existing literature [56, 114], as will be discussed later. To briefly explain, Microsoft’s Live.com search interface (from which the data was collected), is structured as follows:

- For sponsored search there are up to 8 available slots (positions) in which sponsored URL links can be placed. Three of these positions (ranked as 1-3) appear at the top of the page, above the organic search results, but delimited from those by a different background. In addition, the page can display up to 5 additional links in a side bar at the right of the page.
- The “organic” search results are usually returned as 10 URL links/page (a user can opt to change this setting, but very few actually do).

All the sponsored links are allocated based on an auction-like mechanism between the set of interested advertisers (such a display, in any position is called in “impression”). However, the advertisers only pay if their link actually gets clicked - i.e. “pay per click” model. The exact algorithm used by the engine to determine the winners and which advertiser get which position is a complex mechanism design problem and not all details are made public. However, in general, it depends on such factors as the price the bidder is willing to pay per click, the relevance of the query to her set of terms, and her past performance in terms of “clickthrough rate” (i.e. how often links of that user were clicked in the past, for a given keyword). By contrast, in organic search, returned results are ranked simply based on relevance to the user’s query.




Figure 7.2: A (left-side): Cumulative percentage distribution of the number of clicks advertisers in the market receive, wrt. to their rank position, considering the top 5000 advertisers in the market (normal scales). B (right): Log-log scale distributions of the number of impressions, respectively number of clicks, received by the top 10000 advertisers in the market. Note that both distributions follow approximately parallel power laws, but the click distributions levels off in a “long tail” after the first 4000 advertisers, while the impression distribution has a much longer tail (not all appearing in the figure).

7.3.1 Results on display position bias and interpretation

Results for the position bias on click distribution are plotted in Fig. 7.1: part A (left side) for sponsored search and part B (right side) for the organic search. Note that both of these are cumulative distributions: they were obtained by adding the number of clicks for a link in each position, irrespective of the exact context of the queries or links that generated them. Furthermore, both are drawn in the log-log space.

There are two main conclusions to be drawn from these pictures. For the sponsored search results (Fig. 7.1.A). The distribution across the 8 slots seems to resemble a straight line, with a slope parameter approx. $\alpha = 2$. However, such a conclusion would be too simplistic: there is, in fact, a difference between the slope between the first 3 positions (up to $\log_2 3$, on the horizontal axis), and the last 5 positions. The slope for the first 3 positions is around $\alpha_1 = 1.4$, while for the last 5 is around $\alpha_2 = 2.5$. The most likely reason for this drop comes from the way the Live.com search interface is designed. The first 3 slots for sponsored search links are shown on the top of the page, above the organic search results, while the last 5 are shown in a side bar on the right of the page.

Fig. 7.1.B corresponds to the same plot for organic search results, the main effect one notices is the presence of several levels (thresholds), corresponding to clicks on different search pages. We stress that, since this is a log-log plot, the drop in attention between




Figure 7.3: Distribution of advertiser market share, based on their ordered rank vs. the number of clicks their links receive (log-log scales). The left-hand side plot (part A) gives the total number of clicks an advertiser received for all impressions of her links, regardless of the position they were in. The right-hand side (part B) gives the number of clicks received, both in total, but also when her ads were displayed on a specific position on the page (among the 8 ranked slots of the sponsored search interface).

subsequent search pages is indeed very large - about two orders of magnitude (i.e. the top-ranked link on the second search page is, on average, about 65 times less likely to be clicked than the last-ranked link on the first page). The distribution of intra-page clicks, however, at least for the first page of results, could be roughly approximated by a power law of coefficient $\alpha = 1.25$.

All this raises of course the question: what do these distributions mean and what kind of user behaviour could account for the emergence of such distributions in sponsored search results? First, we should point out that the fact that we find power law distributions in this context is not completely surprising. Such distributions have been observed in many web and social phenomena (to give just one example, in collaborative tagging systems, in the work by one of the co-authors of this chapter [93] and others). In fact, any model of “top to bottom” probabilistic attention behaviour, such as a user scanning the list of results from top to bottom and leaving the site with a certain probability by clicking one of them could give rise to such a distribution. Such models and their refinements have been proposed in previous literature [21, 56, 114]. Of course, more fine-grained models of user behaviour are needed to explaining click behaviour in this context. But for now we leave this issue to further research, and we look at the main topic of this chapter which is examining the structure of the sponsored search market itself.

7.4 Market structure at the advertiser level

In this Section, we look at how sponsored search markets are structured, from the perspective of the participants (i.e. advertisers that buy search slots for their URLs). More specifically, we study how relative market shares are distributed across link-based advertisers. We note that in many markets, an often cited rule, also informally attributed to Pareto, is that 20% of participants in a market (e.g. customers in a marketplace) drive 80% of the activity. Here, we call this effect the “market concentration”.

In a sponsored search market, the main “commodity” which produces value for market participants (either advertisers and the search engine) is the number of clicks. Therefore, the first thing that we plotted (first, using normal, i.e. non-logarithmic axes) is the cumulative share of different advertisers (see Fig. 7.2. A. - left side graph). From this graph, one can already see that just the top 500 advertisers get roughly 66% (or about two-thirds) of the total 7.8 million clicks in the available data set³.

Since in our data, there are *at least* 10000 distinct advertisers (most likely, there are many more, but we only considered the top 10000), this means that a percentage of less than 5% of all advertisers have a two-thirds market share. This suggests that sponsored search markets are indeed very concentrated, perhaps even more so than “traditional” real-world markets.

An issue to be discussed here is what this market concentration in terms of received clicks means for the concentration in terms of revenue. While it is user clicks represent the paid “commodity” in most of today’s sponsored search markets, a click may be worth more or less, depending on its position (i.e. rank) on the screen and on how popular (i.e. in demand) are the search keywords leading to that click. Although we have no data for actual monetary revenue, we hypothesize that the market concentration in terms of number of clicks represents a good lower bound for the concentration in terms of revenue as well. This is because the top advertisers are more likely to get their ads displayed in the top positions, and also for the more popular keywords, than advertisers with a low market share. Therefore, they probably pay *per click* at least the average market price, if not more.

7.4.1 Distribution of impressions vs. distribution of clicks for the top advertisers

Next, we studied the detailed distribution of the numbers of impressions (i.e. displayed URLs) and clicks on these impressions, for the top 10000 distinct advertisers. Results are shown in Fig. 7.2.B. (right-hand side graph), using a log-log plot.

The main effect that one can see from Fig. 7.2.B. is that the distribution of impressions and the distribution for clicks received by the advertisers form two approximately parallel,

³Note that an advertiser was taken, following the available data, by the domain URL of the sponsored link. This is a reasonable assumption, in this case. For example, Ebay uses many sponsored links to different products, each relevant for different search terms. However, using this technique, Ebay is taken as one advertiser, regardless of how many different items its URLs point to.

straight lines in the log-log space (i.e. they are two power laws of approximately the same slope coefficient α). There is one important difference, though, which is the size of the “long tail” of the distribution. The distribution of the number of clicks (lower line), levels off after about 4000-5000 positions. Basically, in data terms, this means that advertisers beyond the top 5000 each receive a negligible number of clicks, at least in the dataset we examined. The reason for this may be that their ads almost always appear in the lower display ranks, or simply that they bid on a set of rarely used (or highly specialised) search keywords. By contrast, the distribution of impressions still continues for many more positions (although we only represent the top 10000 distinct advertiser IDs here, as the rest do not play any significant role in the click market).

7.4.2 Distribution of market share per display rank position

The previous Section examined the power law distributions of the number of clicks each advertiser gets *in aggregate* (i.e. over all display ranks his/her links are shown in). Here, we look how an advertiser’s market share distribution is affected when broken down per display rank (an issue we already touched on in Sect. 7.3).

However, we first make a slight restriction in the number of advertisers we consider. As shown in Sect 7.4.1 above, there is a power law distribution in the clicks received by the top 4000 advertisers, advertisers ranked beyond this position each receive a negligible number of clicks. Therefore, in this Section, we restrict our attention to the top 4000 advertisers. As these 4000 advertisers receive over 80% of all 7.8 million clicks in the data set (see Fig. 7.2.A), we do not risk losing much useful information.

Results are shown in Fig. 7.3. First, in Fig. 7.3.A. we show again, more clearly, the power law distribution of the number of clicks for the top 4000 advertisers. Note that this is a “wide” distribution, in the sense that it covers 4000 positions and several orders of magnitude. On the right-hand side graph (Fig. 7.3.B), we show the same graph, but now, for each advertiser, we also break down the number of clicks received by the position his/her sponsored URL was in when it was clicked.

Surprisingly, perhaps, the smooth power law shape is not followed at the level of the display rank - in fact, for the lower levels the variance becomes so great that the distribution breaks down, at the display rank level. We hypothesize the most likely reason for this variance is the way each individual advertiser does the bidding for the preferred keywords at different points in time, or the way he specifies the way his keyword budget could be used in different periods. For example, some advertisers may have a short-running sale campaign, when they will bid aggressively for the preferred keyword, hence getting the top spot. By contrast, others may prefer to have longer-running ads, even if they don’t get the top spot every time. Some anecdotal evidence from online marketing suggests that even just the repeated display of a link of a certain merchant on the screen may count: if a user sees an ad repeatedly in his/her attention space, that may establish the brand as more trustworthy.

In Fig. 7.3.B, by looking the the top 4 advertisers in this dataset, one can already see that

users ranked 2 and 3 utilize a rather different strategy than “the trend” represented by users 1 and 4. While their total number of clicks does follow, approximately the power law, they seem to get, proportionally speaking, more clicks on the top-ranked slot on the page than the rest. While, in order to preserve the privacy of the data, we cannot mention who these companies are, it does seem that users 2 and 3 are actually “aggregators” of advertising demand. By this, we mean online advertising agencies or engines (or automated services offered by the platform itself) that aggregate demand from different advertisers and do the bidding on their behalf. Apparently, this allows them to capture, proportionally, more often the top slot for the required keyword. Unfortunately, however, we cannot investigate this aspect further, since the dataset provided does not contain any information about bidding, budgets or financial information in general.

In the following and last Section of this chapter, we turn our attention to a somewhat different problem: how could we use insights gained from analyzing this query data to provide a bidding decision support for advertisers taking part in a sponsored search market.

7.5 Using click data to derive search term recommendations

The previous Sections of this chapter used complex systems analysis to provide a high-level examination of the dynamics of sponsored search markets. In this Section, we look at how such query log data could be used to output recommendations to individual advertisers. This should lead to answers to questions such as: What kind of keyword combinations look most promising to spend one’s budget on, such as to attract a maximum number of relevant user clicks?

While the previous analysis of power-law formation was done at a macro-level, in this Section we take a more local perspective. That is, we do not consider the set of all possible search terms, but rather a set that is specific to a domain. This is a reasonable model: in practice, most advertisers (which are typically online merchants), are only concerned with a restricted set of keywords which are related to what they are actually trying to sell.

For the analysis in this chapter, we have chosen as a domain 50 keywords related to the tourism industry (i.e. online bookings of tickets, travel packages and such). The reason for this is that much of this activity is already fast moving online (e.g. a very substantial proportion of, for example, flight tickets and hotel reservations are now carried out online). Furthermore - and perhaps more important - there are low barriers of entry and the field is not dominated by one major player. This contrasts, for example, other domains, such as the sale of Ipods and accessories, where Apple Stores can be expected to have a dominant position on the clicks in the market.

7.5.1 Deriving distances from co-occurrence in sponsored click logs

Given a large-scale query log, one of the most useful pieces of information it provides is the co-occurrence of words in different queries. Much previous work has observed that the fact that two search keywords frequently appear together in the same query gives rise to some implicit semantic distance between them [93].

In this chapter, we take a slightly different perspective on this issue, since, in computing the distances, we only use those queries which received at least one sponsored search click for the text ads (i.e. URLs) displayed alongside the results. We argue this is a subtle but very important difference from simply using co-occurrence in organic search logs. The fact that queries containing some combination of query words lead to a click on a sponsored URL implies not only a purely semantic distance between those keywords, more important for an advertiser, the fact that users searching on those combinations of keywords have the possible intention of buying things online.

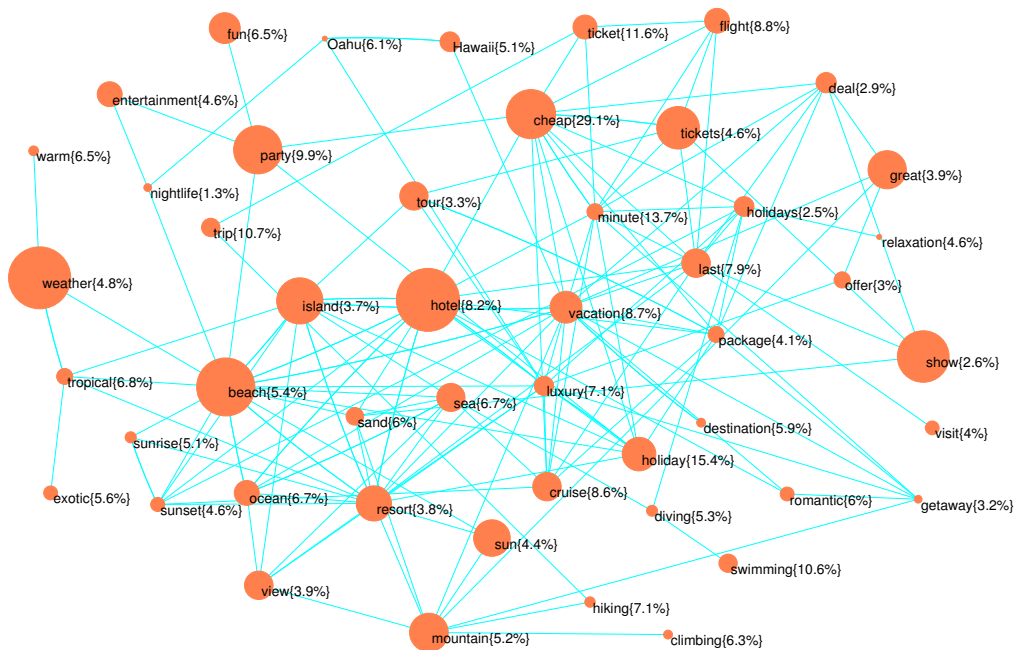
Formally, let $N(T_i, T_j)$ denote the number of times two search terms T_i and T_j appear jointly in the same query, if that query received at least one sponsored search click. Let $N(T_i)$ and $N(T_j)$ denote the same number of queries leading to a click, in which terms T_i , respectively T_j appear in total (i.e. regardless of other terms they co-occur with). Then, the cosine similarity distance between terms T_i and T_j can be defined as:

$$Sim(T_i, T_j) = \frac{N(T_i, T_j)}{\sqrt{N(T_i) * N(T_j)}} \quad (7.2)$$

7.5.2 Constructing keyword correlation graphs

The most intuitive way to represent similarity distances is through a keyword correlation graph. The results from our subset of 50 travel-related terms are shown in Fig. 7.4. In this graph, the size of each node (representing one query term) is proportional to the absolute frequency of the keyword in all queries in the log. The distances between the nodes are proportional to the similarity distance between each pair of terms, computed Eq. 7.2, where the whole graph is drawn according to a so called “spring embedder”-type algorithm. In this type of algorithm, edges can be conceived as “springs”, whose strength is indirectly proportional to their similarity distance, leading to cluster of edges similar to each other to be shown in the same part of the graph.

There are several commercial and academic packages available to draw such complex networks. The one we think is most suitable - and which was used for graph Fig. 7.4 - is Pajek (see [12] for a description). Note that not all edges are considered in the final graph. Even for 50 nodes, there are $\binom{50}{2} = 1225$ possible pairwise similarities (edges), one for each potential keyword pair. Most of these dependencies are, however, spurious (they represent just noise in the data), and our analysis benefits from using only the top fraction,



Note that these click-through rates may, at a first glance, seem on the low side: in general only a few percent of all queries actually lead to a click on an sponsored (i.e. advertiser) link. Nevertheless, as a search engine receives millions of queries in a rather short period of time, even a 5%-10% click-through rate can be quite significant. Note that some keywords (such a “cheap”) have a higher click-through rate than others. The reason for this may be that people searching for “cheap” things (e.g. cheap airline tickets, cheap holiday packages, hotel rooms etc.) may already have the intention to buy something online, and therefore are more likely to [also] click on sponsored links.

However, the most interesting effect to observe in Fig. 7.4 are the term clusters that emerge in different parts of the graph, from the application of the spring-embedder visualization algorithm. For example, the leftmost part of the graph has 4 terms related to weather, such as “warm”, “tropical” and “exotic”. On the top left part of the graph, one can find terms such as “entertainment”, “nightlife”, “party” and “fun”, while very bottom part includes related terms as such “climbing”, “hiking” and “mountain”. The top-right part includes commercial terms such as: “ticket”, “tickets”, “flight”, “cheap”, “last”, “minute”. The central part of the graph includes terms such a “beach”, “sand”, “sea”, “resort”, “ocean”, “island” etc. Additionally, pairs of terms one would naturally associate do indeed appear close together, such as “romantic” and “getaway” and “sunset” and “sunrise” and “ocean”.

In the following, we discuss an algorithm that can detect such clusters automatically. More precisely, we would like an algorithm that selects combinations of tags that look promising in attracting queries and clicks.

7.5.4 Automatic identification of sets of keywords

In this Section, we show how keyword graphs could be automatically partitioned into relevant keyword clusters. The technique we use for this purpose is the so called “community detection” algorithm [164], also inspired by complex systems theory. In network or graph-theoretic terms, a community is defined as a subset of nodes that are connected more strongly to each other than to the rest of the network (i.e. a disjoint cluster). If the network analysed is a social network (i.e. vertexes are people), then “community” has an intuitive interpretation. However, the network-theoretic notion of community detection algorithm is broader, has been successfully applied to domains such as networks of items on Ebay [113], publications on arXiv, food webs [164] etc.

Graph partitioning and community detection algorithm

The basic algorithm used to perform the partition is the same as in the Chapter 6 of this thesis (and adapted from [164]), with the important difference that this time applied to graphs constructed based on click data. We do not provide again the full formal discussion in this chapter, but basically the same type of considerations, regarding computational efficiency, still apply.

Cluster 1	Cl. 2	Cl. 3	Cl. 4	Cl. 5	Cl. 6	Cl. 7	Cl. 8	Cl. 9
beach luxury hotel island resort sun mountain ocean hiking climbing sea sand	party entertainment nightlife fun Hawaii Oahu	package vacation holidays destination deal tour offer great	weather exotic tropical warm	getaway romantic	diving swimming	cruise sunrise sunset	show tickets ticket cheap flight	last minute visit
Keywords eliminated to increase modularity: holiday, holidays, relaxation, trip.								

Figure 7.5: Optimal partition of the set of travel terms in semantic clusters, when the top 150 edges are considered. The partition was obtained by applying Newman’s automated “community detection” algorithm to the graph from Fig 7.4. This partition has a clustering coefficient $Q=0.59$.

7.5.5 Discussion of graph partitioning results

The results from the graph partitioning algorithm, showing the partition maximises the modularity Q for this setting, is shown in Fig. 7.5. Note that this is not the only possible way to partition this graph - if one would consider a different number of strongest dependencies to begin with (in this case we selected the top 150 edges, for 50 keywords), or a different stopping criteria, one may get a somewhat different result. Furthermore, note that some keywords, which were very general and could fit in several clusters (shown below the figure), were pruned in order to improve modularity, through a separate algorithm not shown here.

Still, the partition results shown in Fig. 7.5 match well what our intuition would describe as interesting combinations of search terms, for such a setting. There is one large central cluster, of terms that all have reasonably strong relations to each other, and a set of small, marginal clusters on the side. The large cluster in the middle could be further broken by the partition algorithm, but only if we force some other stop criteria than maximum modularity (such as a certain number of distinct clusters). The partition in Fig. 7.5 fits well with what can be graphically observed in Fig. 7.4: actually, most of the clusters obtained automatically after partition can be identified on different parts of the graph. This does not have to be a one-to-one mapping, however, because in a 2D drawing, the layout of the nodes after “spring embedding” may vary considerably and, furthermore, there are keywords which could fit well into 2 clusters, and were assigned to one as that had a slightly higher modularity.

It is important to stress that the above partition (and resulting vocabulary) was obtained using sponsored search data, rather than organic search one. This is because the purpose of

the resulting vocabulary is to inform the advertisers in this space which keyword combinations the users that actually click on sponsored search ads are more likely to use. This does not necessarily also have to be the most semantically relevant vocabulary (for which one can also use organic search data).

7.6 Discussion

7.6.1 Contribution of the chapter & related work

Our work can be seen as related to several other directions of research. Similar techniques to the ones used in this chapter have been successfully applied to analyse large-scale collaborative tagging systems (Halpin et al.'07 [93]). Other work proposes models to analyse other web phenomena, such as the dynamics of viral marketing Leskovec et. al. [141] or preference networks for Ebay items (Jin et. al. '07 [113]).

The amount of work which is specifically geared to sponsored search auctions, especially empirical studies, has so far been rather limited (probably not least due to lack of extensive datasets in this field). Much of the work that exists looks mostly at user clicking behaviour, and in particular the bias introduced by a link's display rank on clicking behaviour (such as discussed in Sect. 7.3 of this chapter). Prominent examples of this approach are: Craswell et al. '08 [56] and Joachims et al.'05 [114].

Another important direction of work uses existing intuitions about user clicking behaviour to design different allocation mechanisms for this problem - the work of Borgs et. al. '07 [21] is a good example of this approach. By comparison to this work, this direction is much more theoretical and mostly concerned with game-theoretic issues. There has also been recent work that uses Markov models for sponsored search. For example, Hon-sagar and Cherepanov [104] consider the problem of designing optimal bidding strategies for advertisers with budget constraints, participating in multi-unit, multi-slot auctions.

One paper that is very related in scope to ours, since it also provides an empirical study of search engine advertising markets is Ghose and Yang '08 [84]. This work takes, however, a different perspective on this problem, also due to the different type of data the authors had available. By contrast to this chapter, the data that [84] use comes from a single, large-scale advertiser. This means they do get access to more detailed information (including financial one) and can say more about actual bidding behaviour. By comparison, the data available to us for this study does not contain any detailed financial information, but it does allow us to have a much more global level view of the structure of the whole market (from the perspective of the search engine, not just a single advertiser). This provides very important insights about the structure of sponsored search markets.

There exists previous work that has applied similar co-occurrence-based techniques to organic search logs or tagging systems [51, 53, 93]. However, our focus in this chapter is different: we do not aim to merely deduce what is the semantic distance between keywords

in the general sense, but what kind of combinations of keywords are financially interesting for a sponsored search advertiser to bid on. This is the reason why the size of the nodes and distances computed in Fig. 7.4 are built using only queries which lead to an actual click on a sponsored ad. Basically, this is equivalent to filtering only the “opinion” (expressed through queries) of the subset of users that are likely to buy something online, rather than all search engine users. To our knowledge, this is the first approach to use sponsored search click data in this way. A related recent paper, with a somewhat different focus from ours, is Malekian et. al. '08 [150], who propose a method for recommending efficient query rewrites using pay-per-click search advertising data.

7.6.2 Future work

This work, being somewhat preliminary, leaves many aspects open to future research, of which we only mention a few possibilities. One such aspect would be the issue of *externalities*: how the presence of links by competing advertisers influences the clickthrough rates of other bidders. As the competition is basically on customers' attention space, externalities play an important role in the efficacy of sponsored search impressions.

Another very interesting topic would be to study the structure of sponsored search markets (in terms of advertiser market share etc.) not only at the global, macro-level, but at the level of individual sets of keywords. In fact, sponsored search can be seen not only as one market, as a network of markets, since most advertisers are interested in (and bid on) a specific set of keywords related to what they are selling. For example, we could apply our “community detection” algorithm to partition not only sets of search keywords, but also sets of bidders (advertisers) interested in those keywords. This should allow us to derive more in-depth insights into the structure of sponsored search.

7.6.3 Acknowledgements

The authors thank Microsoft Research for their support, in the framework of a ‘Beyond Search’ award. We also wish to thank Nicole Immorlica and Renato Gomes (Northwestern University) for many useful discussions in the preliminary stages of this work.

Part IV

Conclusions

Chapter 8

Conclusions and further work

This thesis has investigated several aspects concerning modeling of complex preferences in agent-mediated electronic markets. The research work can be classified into three main parts: electronic negotiation, sequential auction bidding and web applications, each consisting of 2 chapters. All the problems identified and investigated here are important, open problems, and each of the chapters brings a novel contribution to existing literature in the field.

It is important to note that the work reported in this thesis, although it's inspired by practical, applicable problems, aims to have a contribution to the fundamental research in this field. This is different from applied, industrial case-based research. While we also describe one industrial case study (which also resulted in publications), this is reported in an appendix.

This final chapter aims to give a summary of the main contributions of this thesis to literature of the research presented in each of the chapters. Furthermore, we discuss which directions we identify as most promising to be explored in future research. We address these two issues in separate subsections.

8.1 Overview of the research contributions per chapter

The research presented in each of the chapters of this thesis brings a novel contribution to the existing literature. In the following, we briefly summarize the results and contribution of each chapter.

Chapter 2: This chapter shows, in a bilateral, and partially cooperative negotiation framework, how incomplete preference information could be used to improve multi-attribute negotiation outcomes. We propose a mechanism through which varying amounts of preference information (in the form of attribute utility weights communicated by the negotiation partner) can be used to improve the efficiency of bargaining outcomes. Furthermore, when

preference information is not available, we propose a novel “guessing” mechanism to estimate these utility weights from the bids of the opponent. While it is true that the results are tailored to the negotiation framework for which they were developed [116–118], at the moment the original publication appeared, the results provided an important proof-of-concept application about the use of incomplete preference information in multi-attribute negotiation. In the meantime, the field developed considerably and many authors, either based on or partially inspired by our research, continued to build and develop further models for the multi-issue negotiation problem, in e.g. [25, 78, 103, 109, 136, 201, 225, 233].

Chapter 3: This is a more extensive chapter in the thesis, and provides a substantial contribution to the negotiation literature. The chapter studies the problem of modeling bilateral negotiations over many interdependent, binary issues or over the composition of a large bundle of items. The issue of non-linear preferences had not been much explored in existing literature prior to our research, but was known to be much harder and more computationally expensive. Our work proposed a novel utility graph formalism for modeling agent preferences in such negotiations, assuming the non-linearities can be succinctly represented in k -additive form. We show how such utility graphs can be used for efficient opponent modeling and for focusing the search for an efficient outcome in the most promising part of the complex utility space.

A second contribution is the method proposed to approximate the structure of a utility graph for an anonymous buyer, based on collaborative filtering of past negotiation data. This provides a link between the techniques used to model utilities in agent-mediated negotiation and the techniques used in social recommendation engines. The combination of these two techniques allows bargaining agents to find an efficient agreements relatively fast, even in a high dimensional utility spaces.

Ideas from this work were used as a source of inspiration by other authors working on complex multi-issue negotiation (especially work that considers non-linear utility functions) - see, among others: [49, 76, 100, 102, 136, 137].

Chapter 4: This is the first chapter in the thesis to deal with the complex problem of bidding in sequential auctions. The basic contribution of this work to the literature was to introduce explicit risk profiles when designing bidders’ strategies in sequential auctions. We analyze, for a category of expectations of future auction prices, the effect that a bidder’s risk aversion profile has on her decision-theoretic optimal bidding policy.

Our experimental results show that risk-averse agents have, as expected, less chance of ending up with an incomplete bundle and making a loss in any isolated auction sequence. However, for a longer time horizon, due to the fact that risk-averse agents participate in fewer auction sequences, they make, on average, less expected profit. Furthermore, for some market settings (especially those with multiple types of goods), the fact that bidders are risk averse can also decrease auctioneer revenues.

Chapter 5: The research in this chapter studies the use of priced options to solve the problem of exposure to the risk of a loss that bidders with valuation complementarities face in sequential auctions. This problem, although it often appears in practice, is known to be

hard, in the sense that bidders have no generally dominant bidding strategy to acquire the items they need. Our work builds on and significantly extends an idea first proposed by Juda & Parkes [120], which proposed a mechanism for assigning free options to buyers, under some conditions.

By comparison to [120], in our model, sellers of different items have the choice of either auctioning their items directly, or through a priced option. Each seller sets an exercise price for his/her option and then sell it through an auction protocol. We analyze this model theoretically for the case where the competition is formed by local bidders, and derive the conditions required for both buyers and sellers to have an incentive to use the options mechanism. Furthermore, we also perform an experimental investigation of a market setting in which multiple synergy buyers are active simultaneously.

Chapter 6: This chapter focuses on user preferences in online social systems, in particular tagging systems. It uses data from the social bookmarking site *del.icio.us* to empirically examine the dynamics of collaborative tagging systems and to study how coherent categorization schemes emerge from unsupervised tagging by individual users. First, we study the formation of stable distributions in tagging systems, which are seen as an implicit form of “consensus” reached by the users of the system around the tags that best describe a resource. We show that final tag frequencies for most resources converge to power law distributions and we propose novel methods to examine the dynamics of the process. This convergence analysis is performed both for the most utilized tags at the top, and the so-called “long tail” of tag distributions.

Second, we study the information structures that emerge from collaborative tagging, namely tag correlation graphs. We show how community-based network techniques can be used to extract simple tag vocabularies from the tag correlation graphs by partitioning them into subsets of related tags. Furthermore, we are also able to show, at least for a specialized domain, that shared vocabularies produced by collaborative tagging are richer than the vocabularies which can be extracted from search engine query logs. Ideas from this work were further expanded or served as a partial source of inspiration for a large number of subsequent publications on tagging systems, e.g. [2, 17, 18, 35, 50, 62, 98, 101, 105, 125, 143, 170, 177, 198, 209, 212–214, 234] (among many others).

Chapter 7: This final chapter uses complex systems techniques (similar to Chapter 7) to study the structure and dynamics of online advertising markets. The work presents an empirical study based on real data from Microsoft’s *Live.com*, and it complements the simulation-based study of electronic markets described other chapters. Furthermore, like Chapter 6, this chapter also uses a complex systems analysis perspective.

First, we look at how the display rank of a URL link influences its click frequency, for both sponsored search and organic search. Then, we study the market structure that emerges from these queries, especially the market share distribution of different advertisers. We show that both the number of ad impressions and the number of clicks follow power law distributions of approximately the same slope gradient. However, we find this result does not hold when studying the same distribution of clicks per rank position, which shows considerable variance, due to the way advertisers divide their budget on different keywords.

Finally, we provide a method to represent and visualize keywords of interest in graphical form, as well as a method to partition these graphs to obtain desirable subsets of search terms.

Appendix A: The work in this appendix resulted from a practical case study, which was also used to inform design choices made in other chapters, especially the part that concerns bidding in sequential auctions. It describes an agent-based platform for the allocation of loads in distributed transportation logistics, developed based on an case provided by Vos Logistics Organizing, Nijmegen. The case study around which the simulation was built involves a set of agents bidding for transportation loads to be distributed from a central depot in the Netherlands to different locations across Germany. Our simulation platform supports both human agents, who can bid through specialized planning and bidding interfaces, as well as automated, software agents. The main contribution of this work to the literature is that, unlike other platforms proposed in previous literature to test bidding strategies (such as the ones developed around the TAC competition [192, 230]), it follows a real business scenario proposed by Vos. The distribution of order location and sizes follows real-life data, and the platform includes realistic planning/bidding constraints (such as return or partial truck loads). The fact that the model closely follows a real-life scenario makes it intuitive to use for human transportation planners, and it allows for testing their bidding behaviour against automated strategies that could be designed for this particular industrial setting.

8.2 Further work

In this section, we provide a discussion of the directions of further work we identify as most promising. Some are intuitive extensions of ideas already proposed in different chapters of this thesis, others are more involved and would require a more long-term research effort. For clarity, the discussion is also done per chapter, but paying special attention to the connections that can be identified between the research and techniques presented in different chapters of the thesis.

There are many possible extensions to the bilateral negotiation model presented in **Chapter 2**. As presented here, the model is rather specific, so the conclusions drawn are somewhat tailored to this model. Therefore, it would be interesting to make it more generic, or to study how the heuristics we propose could work (or could be adapted) for more complex, high-dimensional negotiations. Another extension we considered is dealing with other types of incomplete preference information for some attributes, rather than a qualitative ordering heuristic (e.g. for example, if we consider the colour of the car, then fuzzy logic techniques, such as those proposed in [71, 204] could be incorporated). One of the main limitations identified for the negotiation model in Chapter 2, i.e. the assumption of linear utility functions of bargaining agents using the mechanism, has already been extensively investigated by our work in Chapter 3. Furthermore, other very important directions for further work (especially the human-machine interaction aspects) have been researched in subsequent work by the group of my co-author, Catholijn Jonker and collaborators [23–25, 102, 103].

The model proposed **Chapter 3** also raises some interesting ideas for possible extensions. One of them concerns the class of utility graph structures for which our algorithms were designed. The utility graph structures considered in Chapter 3, although we extensively investigated several structures, can be classified as “random graphs” [20] (that means, on average, two random vertexes have the same chance of connected by an edge). This is a standard choice in graph theory. However, recent empirical evidence from large web-based phenomena suggests that real web item graphs (e.g. books on Amazon) may have a more scale free structure. In scale free graphs, the connectivity of nodes follow highly skewed power law distributions (see Chapter 6 of this thesis), where several items are very popular and connected to many others and others much less so. In this case, the separation and learning algorithms discussed in Chapter 3 could be adapted to be made more efficient for scale free structures. Another possibility, inspired by our more empirical work, would be to use the community-based graph identification algorithms discussed in Chapters 7 and 8 to first divide the utility graphs into separate subsets (bundles) of items. Such a prior clustering into subsets of items (which could then be negotiated on independently, in parallel threads) would probably make the negotiation problem more tractable computationally.

Finally, a point of further research which, we argue, is worthwhile to pursue by the research community in the field is coming up with some precise benchmark (or set of benchmarks) for the classes of non-linearity in utility functions which would be useful and practically relevant to study in the context of multi-issue negotiation. There has been, recently, an increasing interest in this field, and different papers propose different models, each with their own techniques and own approach to the problem. Some examples include: simulated annealing, evolutionary techniques, utility graphs, ISO-utility based techniques, even eliminating non-linearities, for some settings [102, 108, 109, 126, 137, 144, 185, 186]. However, the different choices and underlying assumptions that these models start from makes a meaningful comparison difficult. Such a benchmark would not only ease comparison, but would also indicate what kind of settings or utility classes would be interesting to study in more depth for multi-issue negotiation. In this context, empirical work on any available actual e-commerce data from real users could also have an important role to play - as electronic commerce, broadly defined, is the basic “target domain” indicated by most of these techniques.

Chapter 4, which deals with risk aversion, opens up many follow-up research questions, of which Chapter 5 has been one possibility. Especially, in the experimental results, more complex utility functions for the risk-averse bidding agents could be considered, as well as the presence of several complementary-value bidders in the same market. In such a case, there is no dominant bidding strategy in sequential auctions, and the agents would need to dynamically learn and adapt their strategies online, based also on the strategies and risk aversion of the competitors.

Another promising idea would be to combine the option techniques proposed to deal with the exposure problem in sequential auctions from Chapter 5, with risk aversion of the bidders discussed in Chapter 4. This would represent the next, challenging step in addressing the problem of reducing the exposure to risk of loss that agents with complementary valuations

face when participating in sequential auctions. We conjecture that it is quite possible that the benefits of using options would increase in the case agents forming the market have different risk aversion profiles.

Other than the risk averse aspect, the experimental results in **Chapter 5** could also be extended in other ways. For example, while we did consider the presence of several synergy bidders in the market, those experiments still represent only a “proof of concept” regarding the working of the mechanism. In a larger market, especially one in which multiple sellers and buyers interact repeatedly, learning may also play a crucial role. In such a market, buyers could learn their best bidding strategy, but sellers can also learn what exercise price they should fix for their items, such as to obtain maximal expected profit from a sequence of auctions. Because our definition of options is very flexible (e.g. direct sale appears as a particular case, for $K = 0$), sellers of some types of items may conclude that direct auctioning is their best strategy, while the sellers of other items may prefer to use options.

The tagging work reported in **Chapter 6** already underwent several rounds of expansion since our initial result, published after the Santa Fe summer school. But tagging is a very active research area, and several authors have already taken up ideas from our WWW’07 conference paper and pursued them further. Dellschaft & Staab [62], for example, propose a novel generative model to explain the shape of tag power law distributions, by introducing background user knowledge. Heyman, Koutrika & Garcia-Molina [101] perform a more systematic comparison of the dynamics of tagging systems vs. standard web search. Other papers build more elaborate types of graph to examine information from tagging systems. For our own research, one idea we found worth pursuing (if time is available) is more on the social aspects of the problem: examining how sub-communities of users form, and how the local tag vocabularies for these communities emerge (which may be different from the “general” vocabulary of all users). For example, we could use our graph and partitioning algorithms to first partition users in sub-communities, and then examining the tag structures each community uses. A possible step in this direction is the work resulting from our Dagstuhl paper, which looks at the social dynamics of Flickr groups [10].

The work on sponsored search reported in **Chapter 7** is fairly preliminary, although it puts forward some promising ideas. Some aspects we wish to investigate in further work could be, for example, how the keyword vocabularies and graphs extracted from sponsored search click data compare with those extracted from organic search data, or from collaborative tagging. Another important direction is the issue of externalities: how the presence of an ad from a certain advertiser influences the number of clicks that ads of competitors receive, when displayed side by side. Potentially, such externality effects could also be represented in a (directed) graphical form and analysed. Moreover, insights gained from empirical analysis of sponsored search data can also be used to inform research on designing automated bidding strategies (especially as most ad auctions involve repeated and sequential interactions). For example, it could be potentially contribute useful ideas to a new Trading Agent Competition that focuses on ad auctions and sponsored search, proposed to start in 2009 [119].

Finally, there is considerable future work which should be performed regarding the transportation logistics platform described in **Appendix A**. Our work on this case study has now

led to a well-defined platform, which was considered realistic and intuitive for the human logistic planners. However, more applied and behavioural type of research could test their bidding behaviour in sequential auctions, based on our tool. Potentially, such tests could evolve not only human planners, but also intelligent automated bidding heuristics, for which the basis were laid in the more theoretical chapters of this thesis. These ideas probably involve longer term research, which could be done by Vos Logistics itself, or as part of a more applied, industrial research project.

8.3 Concluding remarks

This thesis has investigated some important issues regarding agent-mediated electronic markets. While the topics investigated in the different chapters are rather diverse, they made some solid contributions to the literature, as can be seen from the list of resulting publications and, for some, also from the list of resulting citations. As we see it, all topics covered here relate to the crucial issue of efficiently representing preferences in such complex, online settings - and designing efficient algorithms that can model decision making based on these preferences.

Returning to the opening paragraph of the introduction, it is fair to say, however, that we are still some way off from achieving the vision of fully automated markets, populated by fully autonomous agents taking complex decisions on behalf of their users. Nevertheless, we hope this thesis convinced the reader that, in order to make this vision closer to reality, having good preferences models is a crucial ingredient, and, furthermore, both theoretical, simulation-based and empirical approaches are needed to get us closer to this goal.

Part V

An industrial application case

Appendix A

A Platform for Auction-Based Allocation of Loads in Transportation Logistics

A.1 Introduction

Different chapters of this thesis presented theoretical investigations of the use of automated negotiation and auction mechanisms to allocate resources between self-interested agents. In this appendix, we present the results from a concrete case study performed on the applicability of such techniques to an important practical setting: that of transportation logistics. The work presented in this appendix can be characterized as applied research, i.e. developing an auction platform around a business case study - rather than a fundamental contribution, as other chapters of this thesis. To clearly mark this difference, it is included as an appendix, rather than a separate chapter.

Transportation logistics and supply chain management represents a challenging, but potentially very fruitful area for the application of agent-based electronic market techniques, such as auctions. The increasing complexity and shifting structure of modern supply chains, as well as increasing competitive pressures in this market has led to an increasing demand and interest for such distributed optimization techniques, involving multiple parties. The practical impact of improved allocation which can be achieved through such techniques can be significant. For example, in the Netherlands, the average transport performance is between 40% and 60%. Improving this utilization rate is also the goal of the DEAL (Distributed Engine for Advanced Logistics) project, which groups together several universities and large logistics service providers in the Netherlands. The work reported here (and much of the research leading to this thesis) was also carried out in the framework of this project, involving two of the main partners, namely CWI, Amsterdam and Vos Logistics Organizing,

Nijmegen.

A.1.1 The multi-party logistics domain

Several trends have recently produced a significant impact on the area of transportation logistics. One of these is an increase in competition, with the continual entry of new carriers in the market pushing down expected profit margins. Another one is the increasing complexity and sophistication of modern supply chains. In fact, due to increasing and shifting trade patterns, not only transportation chains have become more dynamic, but also their structure has become increasingly complex.

For example, nowadays it is no longer the case that the company that accepts a transportation order also owns the actual capacity (i.e. trucks) to carry it. Often, multinational companies with large, regular amounts of cargo to be delivered prefer to outsource these orders to other companies that undertake to find convenient delivery options, within a set of pre-negotiated terms. These intermediary logistic companies then negotiate how to distribute these orders with other smaller companies who have the actual transportation capacity (which own the actual trucks and hire the drivers). This can be actually a cheaper option in many cases, as smaller transportation companies often do not have the complex cost structure that larger companies have [221,222].

In standard transportation management literature [221] such distributed supply chains are called multi-party logistics. Existing literature [221] identifies several classes of logistic provider companies, based on the type of services they offer. Although there is some disagreement about the exact usage of the terms, in our approach (and the remainder of this appendix) we use the term 3PL company (third-party logistics providers) to denote those that have their own transport capacity (i.e. truck fleet) and plan this own capacity and 4PL company (i.e. fourth-party logistics provider) to denote those companies which “orchestrate” the supply chain, i.e. acquire large sets of orders from large shippers and then re-distribute these orders among a set of other companies with actual transport capacity.

A.1.2 Company profile

Founded in 1944 as a one-truck company, transporting loads between Oss and Nijmegen in The Netherlands, Vos Logistics has grown into one of the larger logistics service providers in Europe. It has over 3000 trucks, 10000 trailers and containers, 325 storage silos and 2 rail service centers. Vos employs 5000 people working at more than 45 locations throughout Europe, while annual turnover approaches 1 billion euro. The increasing complexity of transportation chains has determined Vos Logistics to offer new solutions to its large corporate customers (shippers), which can now outsource all of their transportation activities to Vos. This lets them avoid the problem of finding and negotiating with individual suppliers, billing, following up orders etc. Another advantage of using this outsourcing service for large shippers is that Vos Logistics has a much better knowledge of the transportation market, so it is better positioned to find suitable sub-contractors. Vos Logistics Organizing from

Nijmegen (henceforth abbreviated VLO in this appendix) is a subsidiary of Vos Logistics B.V. that was set up in order to handle such complex supply chain orchestration activities. Based on the taxonomy above, VLO (the subsidiary) can be seen as a 4PL company, though its parent company, Vos Logistics was founded as a 3PL company and does have its own trucks. Hence, VLO acts as an intermediary company that acquires large (sets of) orders from suppliers and negotiates the allocation of the orders, the terms of transportation (i.e. delivery deadlines, destination) as well as the price at which other carrier companies subcontract these orders.

A.1.3 Automating multi-party logistics using agents

The focus of this work is on automating, through an agent system the second part of the market interaction, i.e. the daily outsourcing of transportation orders to carrier companies who will actually transport them. The first part, which is actually acquiring these orders from large shippers presents less opportunities for automation through a multi-agent system. The reason is that these contracts are usually fewer, larger and closed over a longer time horizon (e.g. a company based in the US may delegate to Vos Logistics Organizing the delivery of the goods imported into Europe over a period of one year). Such large, complex type of decisions cannot be yet expected to be delegated to software agents.

However, allocation of orders on a daily basis to different 3PL carriers was identified as an area with clear potential to benefit from more automated techniques (our previous AAMAS'06 survey paper [222] examined this potential). This automation would involve decision support systems for human planners in the first stage, and next some of the decisions could be delegated to software agents.

A final note is how the allocation occurs in current practice. In the Vos case, negotiation over most orders occurs in a small group of companies who are invited to submit bids for different orders as they arrive in the system. In some cases in which no reasonably priced offer is made, Vos may also solicit other outside companies and carriers to submit a bid (this includes multimodal options, such as rail or water transportation carriers). However, these cases are mostly exceptions (they account for less than 20% of the total value of the orders [222]), so most business is conducted in a group of (up to) 10 companies that can submit bids for a given set of orders. This is the case we are interested in automating through the auction platform presented in this appendix.

A.1.4 Goals of this work

Over the years, several successful auction platforms have been developed in order to allow comparison and evaluation of automated trading strategies to each other. The Trading Agent Competition is, perhaps, the most well known example of this (see [230] for an overview) - most related to this work being its supply-chain version [192]. These platforms are, however, simply not suitable for our basic goal, which is to convince the Vos Logistics Organizing

management (and their partner carrier companies) that agent-mediated electronic auctions can actually be used in practice to automate their daily outsourcing of transportation orders. For this purpose, a custom-based platform was required, modeled around a business case which the planners that actually perform these operations daily can easily recognize and use.

Since the final system is to be used by logistics planners, such a system should closely resemble a real world case, and allow its users to identify the bidding and planning decisions to be taken in this platform as decisions they would usually also take in real life. It should have an interactive, intuitive interface and, moreover, it should seamlessly integrate human agents who take planning and bidding decisions with automated agents implementing an algorithmic strategy or heuristic. This point is especially important for acceptance, since during operational adoption of such a system, it is not realistic to expect that a company would immediately delegate all market decisions to a piece of software, without being confident that such decisions closely model those their human planners would make. To summarize, the goals of this project (and corresponding platform) are:

- The overall goal of the project is to demonstrate the feasibility of applying such an auction system in the day-to-day transportation outsourcing activities of Vos Logistics Organizing (VLO), Nijmegen.
- As a more detailed goal, the platform should allow us to illustrate how different mechanism choices, such as allowing flexible pick-up/delivery times or decommitment [216] (with or without a penalty) can improve efficiency and participant profits.

From an AI or agent researcher's point of view, the developed system can also form a platform to test different aspects of distributed decision making in logistics auctions, more specifically:

- Testing increasingly complex automated trading strategies. At this stage, some very simple strategies have been developed, whose role is mostly to stabilize the market, to make it more realistic. However, more intelligent strategies for this setting can be easily added to the existing platform.
- The demonstrator can also be seen as a platform for analyzing and testing the behaviour of human planners taking part in such an auction.

We wish to emphasize that this appendix is not concerned with proving that any particular bidding strategy, mechanism or scheduling method is superior to others. The readers can consult work which presents and evaluates such strategies, at a more abstract level, in [187, 216, 217]. Rather, our goal in this project is to build an environment which directly models current business practice in transportation logistics (more specifically, a real business case provided by Vos Logistics Organizing, Nijmegen) and in which different analytically-developed strategies can be adapted and tested.

The rest of this appendix is organized as follows. Sect. A.2 provides a high-level overview of our platform and the business case on which it is based. Sect. A.3 describes

in more detail the auctioneer agent, as well as the auction protocol used. Sect. A.4 describes the functionality and behaviour of the automated agents that are currently part of the proposed platform, while Sect A.5 describes the human agent interface and functionality. Sect. A.5 also introduces the cost structure that was used for the agents and the planning assistance interface that was built to assist human planners in taking bidding decisions. Sect. A.6 presents some (very preliminary) results and impressions from a study conducted at Vos Logistics, involving 6 human planners bidding against each other and against our agents, while Sect. A.7 concludes the appendix with a discussion.

A.2 Overview of the business case and our platform

The demonstration takes its starting point in a real-world case of how transportation loads from a depot south of the Netherlands can be distributed across Germany. In order to preserve the privacy of Vos Logistics Organizing, as well as their customers and business partners, some parts of the model are purposely left unspecified or details have been slightly changed, without really affecting how realistic our model is. This especially holds for the names of the customer companies and some specific details about the data used. The main reason for this is that our platform is intended for evaluation not only by planners employed by Vos Logistics, but also by those of some partner companies. The main parts of the problem setting can be summarized as:

- All orders used in the demonstration will be fictive (i.e. randomly generated, not real orders), but, in order to assure the platform is realistic, their destination postcodes, weights, times of delivery etc. are based on real-world distributions.
- All outgoing orders are assumed to be delivered starting from a depot near Maastricht (a town in the south of the Netherlands), while possible return freight (i.e. pick-up) orders appear at destinations across Germany.
- There are n players playing in the role of the carriers (this can vary, we estimated that in our setting it will be up to 10) and one player in the role of VLO (i.e. the auctioneer).
- Each carrier has k trucks to plan (in our demonstration, in order to allow the players to follow all the details simultaneously, we agreed k could be relatively small, e.g. $k = 5..10$). Each truck has a standard capacity of 26 pallets, where pallets are all assumed to have a standard weight of 1000 kg/pallet.

A.2.1 Generating transportation orders

A data set of about 4000 orders was supplied by Vos Logistics, corresponding to orders for a period of time from a real case. These real orders never actually appear in our simulated platform, since that might violate confidentiality agreements between VLO and the shipper

company. However, the orders actually appearing in our platform very closely resemble real orders, as follows.

The German destination (or origin) postcode for each order, which is a two-digit number, was generated as follows. The first digit (corresponding to the broad geographical region), was generated at random using the probability distributions extracted from real data. The weight of the order (expressed in the number of pallets from 1..26), was also generated at random, again from a distribution extracted from the data. In general, some order weights are much more common than others and, furthermore, this also varies by delivery region: some regions receive larger cargo orders, while for some smaller, more frequent orders are the norm ¹. Therefore, the distribution for generating the weight is also dependent on the delivery region (corresponding to the first digit of the postcode). Finally, the second digit of the postcode (which corresponds to a specific town within this general postcode region) was generated at random, but 50% of the weight was given to the 2-3 most important second digits for the area (usually corresponding to a larger town or population center).

In order to have a closed loop demonstration, we assume that the carriers also have return orders available. The return orders are, conceptually, offered by sellers from different areas - although in our demo they will be sold through the same auction mechanism. Outgoing and return orders have asymmetric distributions (60% of all orders are outgoing and only 40% are return orders). This is also realistic for this business scenario, given available data. In real life there are two types of orders: "ON" orders (which must be delivered exactly on their target delivery date) and "BY" orders (which are to be delivered by a certain deadline date, where early delivery is allowed). To simplify the setting, and also allow more competition and flexibility in planning in the simulation, at least for now, all orders in our platform will be considered "BY" orders.

Another very important parameter in such a platform is the lead time of an order, which, roughly defined, represents the difference in days between the time when an order is to be delivered (i.e. the delivery time or deadline) and the time when the order actually appears in the platform (is put up for auction). Here, we also follow a pattern extracted from the real data, as described in the following.

Each order is assigned a random lead-time, produced using a series of adapted, lognormal distributions. The peak of these lognormals will be the first acceptable lead-time day for the order, but with a long tail (see Fig. 1 for an illustration). This means that orders that are to be delivered 3, 4 days or even a week after the minimum lead time can appear, albeit with exponentially decreasing probability. For example, most orders to be generated with a minimum lead time of 1 are to be delivered in the next two days.

The reasons why we need several lognormal distributions is that different types of orders have different lead-times (we identified 3 categories, according to the order data supplied). Thus, orders that are to be delivered to postcode regions in the west of Germany (places closed to the Dutch border) and whose delivery and return trip can be completed within the

¹While we cannot give the full details, a statistically weak, but still significant correlation coefficient of $\eta = 0.4$ was found between the delivery area postcode and the size of an order.

same day have, in general, shorter lead-times than orders that require a minimum of two days travel (including the return trip).

A.2.2 Computing prices and costs

As one would expect in any auction platform, the final price for each order will be determined by the bidding in the open market. However, in an interactive demonstration, we had to build in a mechanism to assure that prices for the orders quickly converge to actual prices (in euro) that human planners would expect to see. Fortunately, also in current practice there is a mechanism to assure this. There is a partner company of Vos Logistics (the name of which, again, we cannot give for privacy reasons), that can transport orders to any destination in Germany. They do provide a standard price scheme which quotes a delivery price for any combination of order size (in number of pallets) and German postcode region. It is very important to stress that these are maximal prices: in general VLO expects to get (and usually gets) much better delivery prices from their closed group negotiation with the partner carriers, otherwise it would be unable to make a profit. The services of this company are only considered if Vos fails to attract a realistic bid for an order from any of the carriers in their closed group (which can sometimes happen, though rather seldom).

However, having such a set of prices is useful in our system, because it provides a benchmark of what kind of prices are realistic. The way we use this information is in designing the bidding strategy of our automated agents, whose bidding strategy will depend on this standard prices (an exact description of the functionality of these automated agents is provided in Sect. A.4). The point of these agents, in this version of the software, is not to beat the human planners, but to assure that the competition bids they see (and implicitly, the bids they have to submit to beat them), are around actual market prices they would encounter in real life. Henceforth in this appendix, we will refer to this set of prices as the standard industry price table.

Finally, a word should be said about cost data. We have also obtained and incorporated in the cost structure of the bidders, detailed information tables about the exact driving times and distances to any postcode location in Germany, as well as realistic estimations of the fixed costs (e.g. driver salaries, truck maintenance) and variable costs per km (including driving tax and fuel costs). These were incorporated in the cost structures of the bidders when planning their routes (a thorough description is provided in Sect. A.5).

A.3 Auction protocol and design of the auctioneer agent

This section describes the main characteristics of the auction protocols used, as well as other characteristics of the auctioneer agent. To allow more planning flexibility, but also to follow current tendering practices, orders with different lead-times are auctioned with slightly different auction protocols, as described below.



Figure A.1: Example distribution of delivery deadlines for orders, in number of days from the present time, for a load with minimum lead time of 1. The date when an order appears in the system corresponds to the origin.

A.3.1 Auction set-up

Loads are auctioned sequentially (or in 3-5 small batches distributed throughout the day). This resembles current transportation practice. Often, loads are offered by different shippers, who have different deadlines throughout the day for placing their orders.

For the current set-up, all auctions are ascending (i.e. English) auctions, but adapted to better fit the actual tendering process, as it is currently performed. There are two main types of auctions, differentiated by their closing protocol.

A.3.2 Auctions for loads with a short lead time

This protocol (more similar to ascending English auctions²), is applied to orders with delivery deadlines which are 1 or 2 days away from the current time. The auction is incrementally descending (lowest offer wins). After the last offer has been placed, the other bidders are given at least 1/2 hour to respond with a new offer, after which the auction closes and the lowest bidder so far is awarded the order. Of course, in our simulated environment 1/2 hour is replaced by 30 seconds to 1 minute. The actual delay to be used (in number of seconds) can be specified by the human user through the interface. Therefore, our auctions have a “soft” closing time (deadline), i.e. they are extended for a short time after the last bid is received, in order to allow other bidders the chance to respond to this bid.

²To be more precise, this extending deadline protocol resembles the most to the protocol used by the e-commerce site Amazon.com.

A.3.3 Auctions for orders with a longer time horizon

For orders with delivery deadlines over 3 days into the future, the simplified protocol cannot be applied, since most bidders do not plan so far in advance. Additionally, some flexibility must be added in the simulations, in order for us to observe the benefits of allowing time window relaxation / the penalty effect for delays.

Therefore, for such orders we use the following decision procedure. For each order, we set a reservation threshold (visible or invisible to the bidders themselves), which gives a reasonable market cost of the order which a shipper would accept in order to have a commitment (without waiting until the last moment to go through the auctions). In our demonstration, the threshold could be set as a percentage below the standard industry price table (as described above) for this configuration of load and destination postcode.

When the order appears in the system, all bidders are informed and can make offers. If a carrier makes a bid that is higher than the reservation price (i.e. not acceptable), then the offer is rejected, the carrier is informed of this and can bid again. A rejected offer (above the reservation price) is thus non-binding to either party, i.e. no commitment exists. If any carrier makes a bid that is below the reservation threshold, and thus acceptable, then all carriers are informed and the auction is moved to the “usual” auction queue (i.e. sold through the auction protocol described in Sect. A.3.2). This means, bidders will have sufficient time to respond after the first offer is made, otherwise the contract is awarded to the initial bidder. If, by 2 days before the deadline, no carrier made a bid in the “acceptable” range (i.e. below the reservation price), then the load is still auctioned using the “usual” procedure, described in Sect. A.3.2.

This protocol ensures that bidders that wish to plan in advance are given the chance to do so, but only if they make a reasonable offer, where by “reasonable” we mean considerably below the price that could be expected to be achieved by waiting closer to the actual deadline. An optional alternative, that could be of interest here, is to allow the human playing the VLO side to change the acceptable reservation threshold during the game, if time passes and an order does not appear to attract enough attention and thus risks remaining undelivered.

Finally, as a future research idea, the reservation threshold could be made dynamic (i.e. automatically increasing), according to a discount function. This function would balance the shipper’s desire of getting a better price for his delivery and the risk of not getting his load delivered in time, as the deadline approaches. This is relatively easy to implement in the current demonstration tool but, at least for the moment, we prefer to focus on testing and usability studies using the simpler setting.

A.3.4 Total capacity of loads to be generated per day

A problem that arises in designing such an agent trading platform is to choose the total capacity of orders which should be generated per day. This choice is an important one, because it gives the player an impression of how “competitive” the whole scenario feels. In

our model, we propose an estimation for this that depends on several parameters:

- n - number of participants representing carriers
- k - number of trucks/participant (our case, e.g. $k = 5$)
- $p = 26$ - number of standard pallets/truck
- s - a coefficient representing the “saturation” of the market. This is an important parameter, which allows us to control the market balance between demand (i.e. coming from outstanding orders) and available supply of transportation capacity.

A rough heuristic evaluation of the capacity of the total capacity of the simulated market we consider will be given by:

$$s * n * k * p$$

Thus, orders will be generated at random using the above distributions, until the total capacity reaches the above value (after choosing the saturation parameter s). This will necessarily be only a very rough estimation: because orders are at random and there are time window constraints, there is no real way to know what is the true capacity of the market - unless we would centrally compute, in advance, the best possible plan for the day for all available trucks. This is not really feasible and it's also not required, because in practice not all capacity of the trucks of a carrier company is allocated in the “closed group” auction. In practice, trucks taking part in such an auction may also acquire loads elsewhere - and they only fill up using the current auction. Furthermore, there should be some differentiation between the capacities of different players.

In order to account for this, we could make the following choice: of the total estimated market capacity, we consider that $i\%$ is filled from other sources (“ i ” stands for the initial fill percentage). Thus, an estimated $s * n * k * p * (1 - \frac{i}{100})$ in total capacity will be filled through the auctions, and $s * n * k * p * \frac{i}{100}$ will be pre-filled, through a heuristic, before the auction starts.

A.3.5 Auctioneer user interface

A screen shot of the auctioneer interface was omitted due to lack of space³, but we provide a brief description of its functionality below. Basically, both the order generation and awarding of orders (i.e. auction closing process) executed by the auctioneer platform can be run in two possible ways:

- *Automated control*: In automatic order generation, the user only specifies the parameters of the generation process (as described above) and the arrival rates of orders in the

³Interface pictures may not be entirely useful, since all interfaces are currently in Dutch, to ease understanding in a business environment.

platform. In automatic tendering mode, the auctioneer waits a number of seconds after the last received bid (which the user specifies through the interface) before making the decision to award the order. This actually varies based on the order lead-time, as described in Sect. A.3.2. Orders with longer lead-times, which remain open for bids until a few days before the delivery deadline, are temporarily shown in a different list and are moved to the “active bidding” queue two days before expected delivery.

- *Human control:* In our interface, a human auctioneer (representing the 4PL company, in this case Vos Logistics Organizing) can make, change or correct any of the decisions taken by the system (either order generation or tendering of orders). We found this is a very useful feature in any live, interactive simulation with several human planners, who first are required to get used to the interface etc. This lets the human auctioneer feel firmly in control of the process, even if he chooses to let the software agent take some of the decisions on his/her behalf.

The switch between these modes can be performed dynamically (and online), by simply checking/unchecking a multi-option box.

A.4 Automated bidders: description and user interface

The role of the automated bidding agents is to ensure the stability of the market and that prices in the demonstrator converge to a realistic level. Therefore, it is enough in a first implementation, if the automated agents use a simple, myopic bidding strategy. The bids are simply based on a standard industry price table (c.f. Sect. A.2.2), which gives a rate for each combination of load/delivery region.

Since this is an English auction, there are two levels, which are randomly determined for each bidding agent: the level of the initial bid and the reservation level (i.e. the lowest the agent will go with his/her bids). Both are generated at random from normal distributions, which are centered at certain levels above and below those taken from our industry price table, as supplied by Vos Logistics. The parameters to be set for automated strategies are:

- Percentage of mean mark-up of the initial bid over the industry price table (and the corresponding dispersion).
- Percentage of the reservation price vs. standard industry price table, for that postcode region and weight (again, this is the mean of the distribution, and a dispersion is also chosen).
- Concession speed (giving how fast the agent’s bids go down from his initial price to the reservation price, i.e. frequency of bidding).
- Number of automated bidders and percentage of orders the automated agents bid on. This gives the pressure that independent bidders apply on the market.

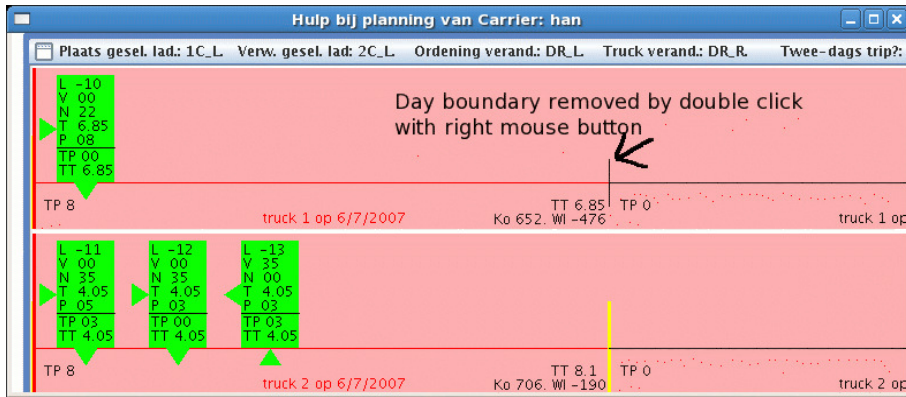


Figure A.2: Basic layout of the planning support window. Each line represents a truck, and each colored container a load (see below for a description of the symbols on each load). For each day, the costs (Ko), profits (Wi) and total traveling times (TT) are computed by the system. Vertical yellow lines represent day boundaries, which can be removed for multi-day planning.

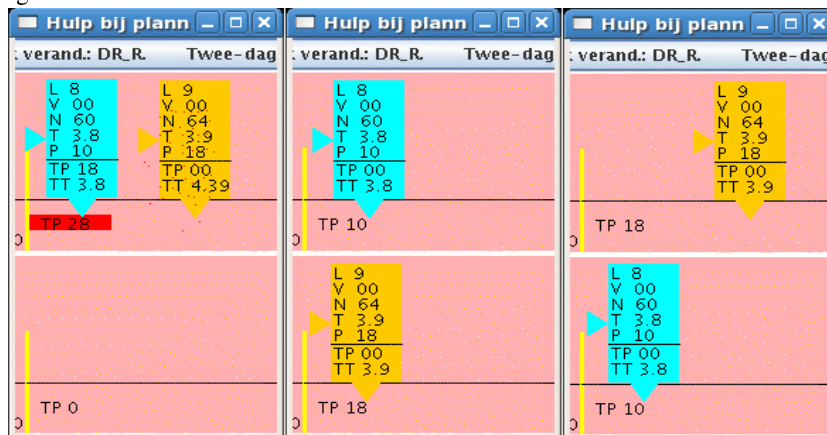


Figure A.3: Left: A number of pallets constraint violation (maximum admitted, 26 pallets/truck), and two possible solutions (center and right), with loads being moved to different days.

A.5 The carrier agents: description and user interfaces

This Section aims to give a technical description of the problem faced by the human carriers in our model and the interface available to them in the demonstrator. More precisely, two distinct interface windows are available to human carriers:

- One for visualizing ongoing auctions for loads and bidding
- One for planning assistance, in which human planners are given a (stylized) impression of their transportation capacity (i.e. trucks) and can visualize and (automatically) determine the way acquired orders fit in their already planned routes, as well as the expected costs.

A.5.1 Transportation model and carrier costs

The transportation planning problem, is, in itself, a complex one to solve. The bidding decisions which the carrier takes are directly dependent on the way the carrier can fit the loads he is bidding on in his/her already existing plans (i.e. how well he/she can form profitable bundles of loads during planning). In turn, this depends on the cost model. Our tool does provide planning assistance, by computing the costs for each combination of loads considered. In our model, costs of each carrier are of two types:

- Fixed costs, per day and per truck. These are expressed as a fixed amount (in euro).
- Variable costs: all these costs are assumed to be proportional to the distance traveled. These are expressed as a cost in euro per kilometer traveled.

Both of these are set to a realistic level, after discussions with Vos Logistics. The distances within Germany, as well as from Maastricht to/from destination postcodes in Germany are computed based on a supplied distance table. This distance table contains, for each pair of first two digits of German postcodes, a distance in kms, as well as a distance in kms from any German postcode to/from Maastricht.

Our planning tool enables the carrier to visualize how filled the trucks are at each time point, the time windows in which loads can be delivered as well as any violation of constraints. There are several types of constraints that need to be met in transportation settings. First there are obvious capacity constraints: a truck cannot be filled at any one time with more than 26 pallets. Second, there is a strict legal constraint about the maximal driving time any driver can actually drive per day - in the EU, this is fixed at 9 hours. Any driving plan has to satisfy these constraints to be feasible.

The tool also provides decision support (see below), by computing the length of the route for the partial daily plan - and, thus, the costs incurred so far, for each possible bid the human planner chooses to make. The length of the route is computed (given the distance table available), through a simple *insertion heuristic*. Insertion heuristics are known to provide a very good approximation of the optimum in small settings - and are known to be computationally more efficient than solving the TSP problem with a more advanced method. Thus, at each point, the expected profit the agent can make so far can also be computed.

A.5.2 Penalty for late deliveries

An issue of relative importance in actual applications is what happens if delivery is (slightly) late, compared to the agreed date⁴. In real life, this does happen to a very small percentage of accepted orders, because profit margins in transportation logistics are tight and carriers have to try to make use of all possible bundling options. Given the business of the underlying customer company, we have decided not to treat slight delays as a strict, inviolable constraint, but to allow orders to be maximum one day late, against payment of a penalty. There are two ways to model the penalty in our system:

- Fixed costs/day of delay (e.g. 50-100 euro for each day the truck is late).
- Proportional, as a percentage of the total value of the transportation order.

In our setting, we currently implement a fixed penalty/day of delay - as opposed to a penalty which is proportional to the value of the order. This is a realistic model, since any delay can be seen as a loss in the reputation of the carrier, regardless of the size or actual value of the order. It is up to the bidding carrier if he chooses to incur this penalty in his planning, but in the current set-up only exceptionally profitable planning configurations would justify the chosen level of penalty for an order. Future versions of the system could consider allowing for differentiated bidding, based on the exact date when the order is delivered (an option discussed in [222]).

A.5.3 Information supplied about other carriers during the competition

An important point to be discussed is what kind of information should be available to human bidders (carriers) in the tool, regarding the activity of the other bidding carriers. This represents a trade-off decision, since on one hand we need to model real life and not compromise the privacy of competing parties, on the other hand in a dynamic simulation environment, agents can be expected to have a reasonable idea about their competition. The following choices have been made:

- Regarding other bids made on existing orders (which the agent is also interested in), the agent should be able to visualize the amounts of the competing bids for the loads he/she is also interested in, but not the identity of the other bidders. Otherwise said, he can see how far he needs to lower his prices to win, but not where the competition for the orders is coming from.
- At the end of each day, a “leader board” is displayed, giving the gross profits rates so far, for all human carriers in the game. We recognize this information about the

⁴As already discussed in Sect A.2, early deliveries are allowed, since we consider all our orders “BY” type of orders.

competition may not be known in real-life, but it may be important in an interactive, game-like simulation scenario for the participants to have a signal of how well they are doing, by comparison to their competition. Also, only knowing the profit margins does not reveal much (if anything) about the bidding strategy and underlying planning of the competing carriers.

A.5.4 Planning and bidding decision support interface

The software developed for human carrier agents has two distinct interfaces: the bidding and the planning support interfaces. In this appendix we only illustrate (in Fig. A.2) some of the features of the planning support interface, as the bidding interface contains relatively straightforward lists of orders one which one can place bids.

The planning interface (see Figs. A.2 and A.3) consists of several horizontal lines, one per each truck that the carrier owns. All trips are assumed to be return trips to/from a depot in Maastricht, for any postcode address in Germany. These trips can be one-day trips, for short-distance orders or two-day trips, for destinations further away (the choice is made by simply clicking a yellow vertical bar).

The interface is a drag-and-drop one, which makes it intuitive and very easy to use. Loads are marked in the system by colored rectangular shapes, marked by two arrows. The side arrows represent pick-up, respectively drop-off points, within the schedule of that day. Each load is marked with: its load no (L), the 2-digit German postcodes of the source (V) and destination(T), number of pallets (P) and time it takes to transport this load (T). The total number of pallets and total traveling time are shown below a black line. Constraint violations will automatically be highlighted in red.

Load symbols can have 3 possible colours:

- *Green*: Loads which have been already acquired (and awarded to the carrier) in auction and which need to be planned for transportation.
- *Light blue*: Loads for which a bid has been placed (thus the agent is bound by the bid he made, since bids are binding), but which have not been won yet by the carrier at the price he offered.
- *Yellow-brown*: Loads which are only placed for tentative planning to see if the planning constraints (total driving time, number of pallets etc.) can still be satisfied given already acquired loads, as well as an estimate of expected profits.

For each truck timeline and day, the system automatically computes the total driving time and the number of pallets loaded and automatically signals (by highlighting in red) if any constraints are being violated. The most useful feature for deciding the minimum bid level is, however, the online computation of the potential profit and loss to be made by inserting a load in the current route. This is basically the difference between the current bid

made for the load and the cost of the extra travel detour for delivery/picking up that load. Empty scheduled already start with a negative profit associated to them, equaling the fixed costs per day and truck.

For loads that have not been bid on yet, but are tentatively dragged & dropped into the schedule, the information about changes in pricing provides very useful information about what is the minimum bid that can be placed if the carrier decides to acquire that load.

A.6 Outline of preliminary human bidding results

A preliminary test of the platform involving 5-6 experienced Vos transportation planners was performed at Vos Logistics. In this test, planners were asked to bid against each other and against our software agents for loads, and their strategies as well as the profit they made with the acquired loads was recorded. Results so far are preliminary, and it was agreed that another large-scale test would be performed in the following months, in order to enable us to extract better empirical data. However, from the testing performed some preliminary conclusions can already be highlighted:

- First, the bidding and planning support interfaces were considered very helpful and realistic by all the planners involved. Some participants even claimed they were superior to the planning system currently being used in everyday planning.
- The presence of automated bidding agents (although they currently only bid based on a randomly perturbed set of industry prices), is crucial for the stability of the market and the convergence of prices to realistic levels.
- The profit levels in the simulation do, very roughly, commensurate with the skill of the bidder. However, in order to ensure that the profit rates actually match current practice, the pricing scheme and other system parameters require some further refinement.
- The planning scenarios considered in the simulation could be expanded to consider some other situations appearing in real life (e.g. multiple one-day return trips).
- Other, more advance functionality could be built into the platform, such as support for combinatorial bidding [194] or allowing the possibility of decommitment for loads already acquired (a possibility analytically studied by us in [216]).

Overall, the planners and managers present were quite impressed with the faithfulness to reality of our platform, and it was agreed that a larger test will be conducted, as well as more concrete steps to be taken towards operational use of such techniques.

A.7 Discussion

Transportation logistics represents an important application area for multi-agent systems, due to its inherently distributed and dynamic nature. Several approaches have been presented in recent years to this problem, some leading to commercially successful, operational systems. The LS/AT system, presented in Dorer & Calisti [64] is one of the most well-known systems that uses agent techniques (mostly constraint-reasoning type techniques) for dynamic transport optimization. The Magenta system [205] is another such system, which explores the use of swarm-based optimization techniques in this setting.

By contrast to these systems, the emphasis in our approach is not directly on optimization of the planning (though that remains, of course, the final goal), but on automating the market interaction between several companies in a multi-party logistics negotiation. Our approach can be seen as creating a testbed, in which each company or carrier can then apply its own optimization and bidding techniques, the performance of these techniques can then be easily measured and compared.

The approach we take is most similar to the work which proposes different trading platforms to test different aspects of bidding and decision making in electronic markets. There are many such platforms proposed in multi-agent literature, the most well-known being the Trading Agent Competition (TAC); the most similar TAC to our approach is, probably, the supply-chain TAC version [192]. Of course, our platform may not have all the sophisticated features of the TAC platforms, but unlike TAC, the starting point of our work was in the applicability of the market setting to a real business case, rather than scientific curiosity or relevance. To the best of our knowledge, it is the first work to describe an agent-mediated auction platform that is modeled around a real-life business scenario, where the orders characteristics, costs, profit margins etc all resemble those encountered in real life.

Another important aspect of our platform is the ability to integrate human bidders and automated trading strategies in the same platform. We feel this is crucial for real business adoption of agent-mediated electronic market techniques because, at least for some of the interacting parties, the human owners will want to remain in control, before delegating any financial decision (e.g. bidding) to a software agent. In multi-agent literature there are some games specifically developed to test human decision-making in negotiation and auctions (a good example is the Colored Trails game [90]), but again our platform has the advantage of allowing us to assess such decisions in a real business environment.

Finally, somewhat related to our approach is work on designing stock market trading platforms to test automated bidding strategies (of which PLAT [124] is a well-known example). While this line of work also uses real financial order data to design a realistic market, the characteristics of stock markets (i.e. double auction setting) is very different from the transportation business case we consider.

We conclude that, overall, our platform did achieve the scope it was built for: to convince Vos Logistics Organizing that the an agent-based approach is a valid solution for their business problem. Nevertheless, there are still many aspects open for further research.

The first would be to conduct a (set of) larger scale experiments to get more detailed human bidding data, and to develop better techniques to analyze this data. The second is to adapt some of the bidding strategies developed analytically in our more theoretical lines of work [187, 216, 217], and test their performance in this environment, both against other strategies and against human planners.

Bibliography

- [1] Osepayshvili A., M.P. Wellman, D.M. Reeves, and J.K. MacKie-Mason. Self-confirming price prediction for bidding in simultaneous ascending auctions. In *Proc. of the 21st Conf. on Uncertainty in AI (UAI-05)*, 2005.
- [2] Fabian Abel. The benefit of additional semantics in folksonomy systems. In *PIKM '08: Proceeding of the 2nd PhD workshop on Information and knowledge management*, pages 49–56. ACM, 2008.
- [3] Chris Anderson. *The Long Tail: How Endless Choice Is Creating Unlimited Demand*. Random House Business Books, 2006.
- [4] Martin R. Andersson and Tuomas W. Sandholm. Leveled commitment contracts with myopic and strategic agents. *Journal of Economic Dynamics and Control*, 25(3-4):615–640, March 2001.
- [5] K.J. Arrow. *Aspects of the Theory of Risk-Bearing*. Y. Hahnsson Foundation, Helsinki, 1965.
- [6] Brian Arthur, John H. Holland, Blake LeBaron, Richard Palmer, and Paul Taylor. Asset pricing under endogenous expectation in an artificial stock market. In *Santa Fe Institute Working Paper 96-12-093*, 1996.
- [7] Robert Axelrod. *The evolution of cooperation*. Elsevier, 1985.
- [8] A. Babanov, J. Collins, and M. Gini. Harnessing the search for rational bid schedules with stochastic search. In *Proc. of AAMAS'04, New York, USA*, pages 355–368, 2004.
- [9] F. Bacchus and A. Grove. Graphical models for preference and utility. In *Uncertainty in Artificial Intelligence UAI-95*, pages 3–10, 1995.
- [10] A. Baldassarri, A. Barrat, A. Capocci, H. Halpin, U. Lehner, J. Ramasco, V. Robu, and D. Taraborelli. The berners-lee hypothesis: Power laws and group structure in flickr. In *Dagstuhl Seminar on Social Web Communities*, Seminar Proceedings 08391. DROPS Server, Schloss Dagstuhl, 2008.

- [11] Yaneer Bar-Yam. The dynamics of complex systems (studies in nonlinearity). *Westview Press*, 2003.
- [12] V. Batagelj and A. Mrvar. Pajek - A program for large network analysis. *Connections*, 21:47–57, 1998.
- [13] S. Bateman, C. Brooks, G. McCalla, and P. Brusilovsky. Applying collaborative tagging to e-learning. In *WWW'07 Workshop on Tagging and Metadata for Social Information Organization*, 2007.
- [14] W. Benn, O. Goerlitz, and R. Neubert. Enabling integrative negotiations by adaptive software agents. In *Cooperative Information Agents III*, pages 335–346. Springer Lecture Notes in AI, vol. 1652, 1999.
- [15] Bettina Berendt and Christoph Hanser. Tags are not metadata, but “just more content” – to some people. In Nicolas Nicolov, Natalie Glance, Eytan Adar, Matthew Hurst, Mark Liberman, James H. Martin, and Franco Salvetti, editors, *Proceedings of the 1st International Conference on Weblogs and Social Media*, Boulder, CO, U.S.A., March 26-28 2007. <http://www.icwsm.org>.
- [16] Tim Berners-Lee, James Hendler, and Ora Lassila. The semantic web. *Scientific American*, 5, 2001.
- [17] Claudio Biancalana, Alessandro Micarelli, and Claudio Squarcella. Nereau: a social approach to query expansion. In *WIDM '08: Proceeding of the 10th ACM workshop on Web information and data management*, pages 95–102. ACM, 2008.
- [18] Kerstin Bischoff, Claudiu S. Firan, Wolfgang Nejdl, and Raluca Paiu. Can all tags be used for search? In *CIKM '08: Proceeding of the 17th ACM conference on Information and knowledge management*, pages 193–202. ACM, 2008.
- [19] S. M. Bohte, E. Gerding, and J.A. La Poutré. Market-based recommendation: Agents that compete for consumer attention. *ACM Trans. Internet Technol.*, 4(4):420–448, 2004.
- [20] Bela Bollobas. *Random Graphs*. Academic Press, London, England, 1985.
- [21] Christian Borgs, Jennifer Chayes, Nicole Immorlica, Kamal Jain, Omid Etesami, and Mohammad Mahdian. Dynamics of bid optimization in online advertisement auctions. In *WWW '07: Proceedings of the 16th international conference on World Wide Web*, pages 531–540. ACM Press, 2007.
- [22] T. Bosse and C.M. Jonker. Human vs. computer behaviour in multi-issue negotiation. In Ito, T., Hattori, H., Matsuo, T., and Zhang, M. (eds.), *Proc. of the First Int. Workshop on Rational, Robust, and Secure Negotiations (RRS'05)*, pages 10–25. IEEE Press, 2005.

- [23] T. Bosse, C.M. Jonker, and J. Treur. Experiments in human multi-issue negotiation: Analysis and support. In *Third Int. Joint Conference on Aut. Agents and Multi-Agent Systems (AAMAS'04)*, pages 672–679. ACM Press, 2004.
- [24] Tibor Bosse, Catholijn Jonker, Lourens van der Meij, Valentin Robu, and Jan Treur. *A System for Analysis of Multi-Issue Negotiation*, pages 253–279. Whitestein Series in Software Agent Technologies. Birkhauser, Springer, 2005.
- [25] Tibor Bosse, Catholijn M. Jonker, Lourens van der Meij, and Jan Treur. Automated formal analysis of human multi-issue negotiation processes. *Multiagent Grid Systems*, 4(2):213–233, 2008.
- [26] Craig Boutilier, Rajarshi Das, Jeffrey O. Kephart, Gerald Tesauro, and William E. Walsh. Cooperative negotiation in autonomic systems using incremental utility elicitation. In *Proc. of the Nineteenth Annual Conference on Uncertainty in Artificial Intelligence (UAI-03)*, Acapulco, pages 89–7, 2003.
- [27] Craig Boutilier, Moisés Goldszmidt, and Bikash Sabata. Sequential auctions for the allocation of resources with complementarities. In *Proc. of IJCAI'99*, pages 527–523, San Francisco, 1999. Morgan Kaufmann.
- [28] Oisín Boydell and Barry Smyth. Capturing community search expertise for personalized web search using snippet-indexes. In *Proc. of the Int. Conference on Information and Knowledge Management (CIKM'06)*, pages 1313–1314. ACM Press, 2006.
- [29] Oisín Boydell and Barry Smyth. From social bookmarking to social summarization: an experiment in community-based summary generation. *Proc. of the 2007 Int. Conf. on Intelligent User Interfaces*, pages 42–51, 2007.
- [30] R. I. Brafman, C. Domshlak, and S. E. Shimony. On graphical modeling of preference and importance. *Journal of Artificial Intelligence Research (JAIR)*, 25:389–424, 2006.
- [31] F.M.T. Brazier, C.M. Jonker, and J. Treur. Compositional design and reuse of a generic agent model. *Applied Artificial Intelligence Journal*, 14:491–538, 2000.
- [32] D. Braziunas and C. Boutilier. Local utility elicitation in gai models. In *Proc. of the Twenty-first Conference on Uncertainty in Artificial Intelligence (UAI-05)*, pages 42–49, 2005.
- [33] Jonathan Bredin and David C. Parkes. Models for truthful online double auctions. In *Proceedings of the 21st Conference on Uncertainty in Artificial Intelligence (UAI'2005)*, pages 50–59, 2005.
- [34] Jakub Brzostowski and Ryszard Kowalczyk. Experimental evaluation of possibilistic mechanism for negotiation partners selection. In *Proc. of Second Int. Workshop on Rational, Robust, and Secure Negotiations in Multi-Agent Systems, Hakodate, Japan*. Springer Lecture Notes in Computational Intelligence (to appear), 2006.

- [35] Adriana Budura, Sebastian Michel, Philippe Cudré-Mauroux, and Karl Aberer. To tag or not to tag -: harvesting adjacent metadata in large-scale tagging systems. In *SIGIR '08: Proceedings of the 31st annual international ACM SIGIR conference on Research and development in information retrieval*, pages 733–734. ACM, 2008.
- [36] Stuart Butterfield. Folksonomy, 2004. <http://www.sylloge.com/personal/2004/08/folksonomy-social-classification-great.html>.
- [37] Peter Buzing, A.E. Eiben, Martijn Schut, and Tudor Toma. Cooperation and communication in evolving artificial societies. In *IEEE Congress on Evolutionary Computation (CEC'04)*. IEEE, 2004.
- [38] Mark M. Bykowsky, Robert J. Cull, and John O. Ledyard. Mutually destructive bidding: The fcc auction design problem. *Journal of Regulatory Economics*, 17(3):205–228, May 2000.
- [39] P. Poupart C. Boutilier, R. Patrascu and D. Schuurmans. Regret-based utility elicitation in constraint-based decision problems. In *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence (IJCAI-05)*, pages 929–934, 2005.
- [40] R. Ferrer Cancho and R. V. Sole. Least effort and the origins of scaling in human language. *Procs. Natl. Acad. Sci. USA*, 100:788–791, 2003.
- [41] Tom Carter. A short trip through entropy to power laws, 2007. Complex Systems Summer School, Santa Fe Institute, NM.
- [42] Ciro Cattuto, Vittorio Loreto, and Luciano Pietronero. Semiotic dynamics and collaborative tagging, 2006. preprint.
- [43] Ciro Cattuto, Vittorio Loreto, and Luciano Pietronero. Semiotic dynamics and collaborative tagging. *Proceedings of the National Academy of Sciences (PNAS)*, 104(5):1461–1464, 2007.
- [44] U. Chajewska and D. Koller. Utilities as random variables: Density estimation and structure discovery. In *Proceedings of sixteenth Annual Conference on Uncertainty in Artificial Intelligence UAI-00*, pages 63–71, 2000.
- [45] Kay-Yut Chen, Leslie R. Fine, and Bernardo Huberman. Predicting the future. *Information Systems Frontiers*, 5(1):47–61, 2003.
- [46] Yann Chevalere, Paul E. Dunne, Ulle Endriss, Jerome Lang, Michel Lematre, Nicolas Maudet, Julian Padget, Steve Phelps, Juan A. Rodriguez-Aguilar, , and Paulo Sousa. Issues in multiagent resource allocation. *Informatica*, 30(3–31), 2006.
- [47] Yann Chevalere, Ulle Endriss, Jerome Lang, and Nicolas Maudet. Negotiating over small bundles of resources. In *Proc. of the 4th Int. Joint Conf. on Auton. Agents and Multiagent Syst. (AAMAS)*, pages 296–302, 2005.

- [48] Yann Chevaleyre, Ulle Endriss, and Nicolas Maudet. On maximal classes of utility functions for efficient one-to-one negotiation. In *Proc. of the 19th Int. Joint Conference on Artificial Intelligence (IJCAI-2005)*, pages 941–946, 2005.
- [49] Yann Chevaleyre, Ulle Endriss, and Nicolas Maudet. Tractable negotiation in tree-structured domains. In *Proc. of the 5th Int. Joint Conf. on Autonomous Agents and Multiagent Systems (AAMAS-2006)*. ACM Press, 2006.
- [50] Ed H. Chi and Todd Mytkowicz. Understanding the efficiency of social tagging systems using information theory. In *HT '08: Proceedings of the nineteenth ACM conference on Hypertext and hypermedia*, pages 81–88. ACM, 2008.
- [51] Steven Chien and Nicole Immorlica. Semantic similarity between search engine queries using temporal correlation. In *WWW '05: 14th Int. Conf. World Wide Web*, pages 2–11, New York, NY, USA, 2005. ACM.
- [52] P.A. Chirita, S. Costache, S. Handschuh, and W. Nejdl. P-tag: Large scale automated generation of personalised annotation tags for the web. In *Proceeding of the 16th Int. World Wide Web Conference (WWW 2007)*, pages 845–854. ACM Press, 2007.
- [53] R. L. Cilibrasi and P. M. B. Vitanyi. The google similarity distance. *IEEE Transactions on Knowledge and Data Engineering*, 19(3):370–383, 2007.
- [54] Dave Cliff. Explorations in evolutionary design of online auction market mechanisms. *Electronic Commerce Research and Applications*, 2(2):162 – 175, 2003.
- [55] V. Conitzer, T. Sandholm, and P. Santi. Combinatorial auctions with k-wise dependent valuations. In *Proc. of the National Conference on Artificial Intelligence (AAAI)*, 2005.
- [56] Nick Craswell, Onno Zoeter, Michael Taylor, and Bill Ramsey. An experimental comparison of click position-bias models. In *WSDM '08: Proc. Int. Conf. on Web search and data mining*, pages 87–94. ACM Press, 2008.
- [57] J. Dang and Huhns M. H. Concurrent multiple-issue negotiation for internet-based services. *IEEE Internet Computing (to appear)*, 2006.
- [58] Rajdeep K. Dash, Nicholas R. Jennings, and David C. Parkes. Computational mechanism design: A call to arms. *IEEE Intelligent Systems*, pages 40–47, 2003.
- [59] Esther David, Alex Rogers, Nicholas R. Jennings, Jeremy Schiff, Sarit Kraus, and Michael H. Rothkopf. Optimal design of english auctions with discrete bid levels. *ACM Trans. Interet Technol.*, 7(2):12, 2007.
- [60] J. K. Debenham. Bargaining with information. In *3rd Int. Conf. on Autonomous Agents & Multi Agent Systems (AAMAS), New York, July 19-23, 2004*, pages 663–670, 2004.

- [61] R. Dechter. *Constraint Processing*. Morgan Kaufmann Publishers, San Francisco, USA, 2003.
- [62] K. Dellschaft and S. Staab. An epistemic dynamic model for tagging systems. In *HT '08: Proc. of 19th ACM Conf. on Hypertext and Hypermedia*, pages 71–80. ACM, 2008.
- [63] Klaas Dellschaft and Steffen Staab. An epistemic dynamic model for tagging systems. In *Proceedings of the 19th ACM Conference on Hypertext and Hypermedia (HYPERTEXT'08)*, pages 71–80. ACM Press, 2008.
- [64] Klaus Dorer and Monique Calisti. An adaptive solution to dynamic transport optimization. In *Proceedings of AAMAS'05 (Industry Track), Utrecht, The Netherlands*, pages 45–51. ACM Press, 2005.
- [65] M. Dubinko, R. Kumar, J. Magnani, J. Novak, P. Raghvan, and A. Tomkins. P-tag: Large scale automated generation of personalised annotation tags for the web. In *Proceeding of the 15th Int. World Wide Web Conference (WWW 2006)*, pages 193–202. ACM Press, 2006.
- [66] P. E. Dunne, M. Wooldridge, and M. Laurence. The complexity of contract negotiation. *Artificial Intelligence*, 164(1-2):23–46, 2005.
- [67] H. Ehtamo and R.P. Hamalainen. Interactive multiple-criteria methods for reaching pareto optimal agreements in negotiations. *Group Decision and Negotiation*, 10:475–491, 2001.
- [68] J Epstein and R. Axtell. *Growing Artificial Societies: Social Science From The Bottom Up*. Brookings Institute Press, 1996.
- [69] Guy Even, Joseph Naor, Satish Rao, and Baruch Schieber. Fast approximate graph partitioning algorithms. *Siam Journal of Computing*, 28(6):2187–2214, 1999.
- [70] P Faratin, C Sierra, and N.R Jennings. Negotiation decision functions for autonomous agents. *Robotics and Autonomous Systems*, 24(3–4):159–182, 1998.
- [71] Peyman Faratin, Carles Sierra, and Nicholas R. Jennings. Using similarity criteria to make issue trade-offs in automated negotiations. *Journal of Artificial Intelligence*, 142(2):205–237, 2003.
- [72] S. S. Fatima, M. Wooldridge, and N. R. Jennings. An agenda-based framework for multi-issue negotiation. *Artificial Intelligence*, 152(1):1–45, 2004.
- [73] Shaheen Fatima, Michael Woolridge, and Nicholas R. Jennings. Optimal negotiation of multiple issues in incomplete information settings. In *3rd Int. Conf. on Autonomous Agents & Multi Agent Systems (AAMAS), New York*, pages 1080–1087, 2004.

- [74] Eric Friedman and David C. Parkes. Pricing WiFi at Starbucks— Issues in online mechanism design. In *Fourth ACM Conf. on Electronic Commerce (EC'03)*, pages 240–241, 2003.
- [75] Katsuhide Fujita and Takayuki Ito. A preliminary analysis of computational complexity of the threshold adjusting mechanism in multi-issue negotiations. In *Proc. of the 3rd Int. Workshop on Rational, Robust and Secure Negotiations in Multi-Agent Systems (RRS'07)*, Silicon Valley, US. IEEE Press (to appear), 2007.
- [76] Katsuhide Fujita, Takayuki Ito, and Mark Klein. Preliminary result on secure protocols for multiple issue negotiation problems. In *Intelligent Agents and Multi-Agent Systems*, pages 161–172. Springer LNCS vol. 5357, 2008.
- [77] N. Garg, H. Saran, and V. Vazirani. Finding separator cuts in planar graphs within twice the optimal. In *Proc. 35th Ann. IEEE Symp. on Foundations of Comput. Sci.*, pages 14–23. IEEE Computer Society, 1994.
- [78] Nicola Gatti, Francesco Di Giunta, and Stefano Marino. Alternating-offers bargaining with one-sided uncertain deadlines: an efficient algorithm. *Artif. Intell.*, 172(8-9):1119–1157, 2008.
- [79] E. Gerding, D.J.A. Somefun, and J. A. La Poutré. Multi-attribute bilateral bargaining in a one-to-many setting. In *Proc. of the AMEC VI Workshop, New York, USA*, 2004.
- [80] E.H. Gerding, D.J.A. Somefun, and J.A. La Poutré. Efficient methods for automated multi-issue negotiation: Negotiating over a two-part tariff. *International Journal of Intelligent Systems*, 21(1):99–119, 2006.
- [81] Enrico Gerding, Rajdeep Dash, Andrew Bye, and Nick Jennings. Optimal strategies for bidding agents participating in simultaneous vickrey auctions with perfect substitutes. *Journal of Artificial Intelligence Research*, 32:939–982, Aug. 2008.
- [82] Enrico Gerding, David van Bragt, and Han La Poutré. Scientific approaches and techniques for negotiation: A game theoretic and artificial intelligence perspective. *Technical Report SEN-R0005*, CWI, Amsterdam, 2002.
- [83] Enrico H. Gerding, Rajdeep K. Dash, David C.K. Yuen, and Nicholas R. Jennings. Bidding optimally in concurrent second-price auctions of perfectly substitutable goods. In *Proc. 6th Int. Joint Conf. on Autonomous Agents and Multi-agent Systems (AAMAS'07)*, Honolulu, Hawaii, pages 267–274, 2007.
- [84] Anindya Ghose and Sha Yang. Analyzing search engine advertising: firm behavior and cross-selling in electronic markets. In *WWW '08: Proc. of the 17th Int. Conf. on World Wide Web*, pages 219–226. ACM Press, 2008.
- [85] Risto Gligorov, Zharko Aleksovski, Warner ten Cate, and Frank van Harmelen. Using google distance to weight approximate ontology matches. In *Proc. of 16th Int. World Wide Web Conference (WWW'07)*, pages 767–775. ACM Press, 2008.

- [86] Scott Golder and Bernardo Huberman. Usage patterns of collaborative tagging systems. *Journal of Information Science*, 32(2):198–208, 2006.
- [87] Ram Gopal, Steven Thompson, Y. Alex Tung, and Andrew B. Whinston. Managing risks in multiple online auctions: An *Options* approach. *Decision Sciences*, 36(3), 2005.
- [88] W. H. Green. *Econometric Analysis*. Prentice Hall, New Jersey, 1993.
- [89] A. Greenwald and J. Boyan. Bidding under uncertainty: Theory and experiments. In *Proceedings of the 20th Conf. on Uncertainty in AI (UAI-04)*, pages 209–216, 2004.
- [90] B. Grosz, S. Kraus, S. Talman, B. Stossel, and M. Halvin. The influence of social dependencies on decision making: investigations with a new game. In *Proceedings of AAMAS'04, New York, USA*, pages 782–791. ACM Press, 2004.
- [91] R. Gutman and P. Maes. Cooperative vs. competitive multi-agent negotiation in retail electronic commerce. In *Proc. Of. the Secod International Workshop on Cooperative Information Agents (CIA'98), Paris*, 1998.
- [92] Mohammad Hajiaghayi, Robert Kleinberg, Mohammad Mahdian, and David Parkes. Online auctions with re-usable goods. In *Proc. 6th ACM Conf. on Electronic Commerce (EC'05)*, pages 165–174, 2005.
- [93] Harry Halpin, Valentin Robu, and Hana Shepherd. The complex dynamics of collaborative tagging. In *Proc. of the 16th Int. World Wide Web Conference (WWW'07)*, pages 211–220. ACM Press, 2007.
- [94] Martin Halvey and Mark T. Keane. An assesment of tag presentation techniques. In *Proceedings of the 16th Int. World Wide Web Conference (WWW 2007)*, pages 1313–1314. ACM Press, 2007.
- [95] Conor Hayes and Paolo Avesani. Using tags and clustering to identify topic-relevant blogs. In Nicolas Nicolov, Natalie Glance, Eytan Adar, Matthew Hurst, Mark Liberman, James H. Martin, and Franco Salvetti, editors, *Proceedings of the 1st International Conference on Weblogs and Social Media*, Boulder, CO, U.S.A., March 26-28 2007. <http://www.icwsm.org>.
- [96] Minghua He, Alex Rogers, Xudong Luo, and N. R. Jennings. Designing a successful trading agent for supply chain management. In *5th Int. Joint Conf. on Aut. Agents and Multiagent Systems (AAMAS-06)*, pages 1159–1166, 2006.
- [97] Marti A. Hearst and Daniela Rosner. Tag clouds: Data analysis tools or social signaller? In *Proc. of 41st Hawaii Int. Conference on System Sciences*. IEEE, 2008.
- [98] Markus Heckner, Tanja Neubauer, and Christian Wolff. Tree, funny, toread, google: what are tags supposed to achieve? a comparative analysis of user keywords for different digital resource types. In *SSM '08: Proceeding of the 2008 ACM workshop on Search in social media*, pages 3–10. ACM, 2008.

- [99] M. Hemaissia, A.E.F. Seghrouchni, C. Labreuche, and J. Mattioli. Cooperation-based multilateral multi-issue negotiation for crisis management. In *Proc. of Second Int. Workshop on Rational, Robust, and Secure Negotiations in Multi-Agent Systems, Hakodate, Japan*. Springer LNCI (to appear), 2006.
- [100] Miniar Hemaissia, Amal El Fallah Seghrouchni, Christophe Labreuche, and Juliette Mattioli. A multilateral multi-issue negotiation protocol. In *AAMAS '07: Proceedings of the 6th international joint conference on Autonomous agents and multiagent systems*, pages 1–8, New York, NY, USA, 2007. ACM.
- [101] Paul Heymann, Georgia Koutrika, and Hector Garcia-Molina. Can social bookmarking improve search? In *Proc. of Int. Conference on Web Search and Data Mining (WSDM'08)*, pages 195–205. ACM Press, 2008.
- [102] K. Hindriks, C. Jonker, and D. Tykhonov. Reducing complexity of an agent's utility space for negotiating interdependent issues. In *International Conference on Complex Systems (ICCS'06), Boston*, 2006.
- [103] Koen Hindriks and Dmytro Tykhonov. Opponent modelling in automated multi-issue negotiation using bayesian learning. In *AAMAS '08: Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems*, pages 331–338. International Foundation for Autonomous Agents and Multiagent Systems, 2008.
- [104] Kartik Hosanagar and Vadim Cherepanov. Optimal bidding in stochastic budget constrained slot auctions. In *EC '08: Proc. of the 9th ACM Conf. on Electronic commerce*, pages 20–20, 2008.
- [105] Ming-Hung Hsu and Hsin-Hsi Chen. A method to predict social annotations. In *CIKM '08: Proceeding of the 17th ACM conference on Information and knowledge management*, pages 1375–1376. ACM, 2008.
- [106] <http://labs.google.com/sets>. Google sets, 2008. Retrieved: 1st Sept.
- [107] John C. Hull. *Options, Futures, and Other Derivatives*. Prentice Hall, fifth edition, 2003.
- [108] Takayuki Ito, Hiromitsu Hattori, and Mark Klein. A multi-issue negotiation protocol for agents : Exploring nonlinear utility spaces. In *Proc. of the Twentieth International Joint Conf. on Artificial Intelligence (IJCAI'07)*, pages 1347–1352, 2007.
- [109] Takayuki Ito and Mark Klein. A multi-issue negotiation protocol among competitive agents and its extension to a nonlinear utility negotiation protocol. In *AAMAS '06: Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems*, pages 435–437, New York, NY, USA, 2006. ACM.
- [110] E. Jacob. Classification and categorization: A difference that makes a difference. *Library Trends*, 52(3):515–540, 2004.

- [111] N.R. Jennings, P. Faratin, A.R. Lomuscio, S. Parsons, M. Woolridge, and C. Sierra. Automated negotiation: Prospects, methods and challenges. *Group Decision and Negotiation*, 10:199–215, 2001.
- [112] A.X. Jiang and K. Leyton-Brown. Bidding agents for online auctions with hidden bids. *Machine Learning*, 67(1-2):117–143, 2007.
- [113] R. Kang-Xing Jin, David C. Parkes, and Patrick J. Wolfe. Analysis of bidding networks in eBay: Aggregate preference identification through community detection. In *Proc. AAAI Workshop on Plan, Activity and Intent Recognition*, 2007.
- [114] Thorsten Joachims, Laura Granka, Bing Pan, Helene Hembrooke, and Geri Gay. Accurately interpreting clickthrough data as implicit feedback. In *SIGIR '05: Proc. 28th Int. ACM SIGIR Conf.*, pages 154–161. ACM, 2005.
- [115] Catholijn Jonker and Valentin Robu. Automated multi-attribute negotiation with efficient use of incomplete preference information. In *3rd Int. Conf. on Autonomous Agents & Multi Agent Systems (AAMAS)*, New York, pages 1056–1063, 2004.
- [116] Catholijn Jonker, Valentin Robu, and Jan Treur. An agent architecture for multi-attribute negotiation using incomplete preference information. *Journal of Autonomous Agents and Multi-Agent Systems (JAAMAS)*, 15:221–225, 2007.
- [117] Catholijn M. Jonker and Jan Treur. A re-usable broker agent architecture with dynamic maintenance capabilities. In *Proceedings of the Third Annual Conference on Autonomous Agents (Agents'99)*, pages 376–377. ACM Press, 1999.
- [118] Catholijn M. Jonker and Jan Treur. An agent architecture for multi-attribute negotiation. In *Proc. of the 17th Int. Joint Conference on AI (IJCAI'01)*, pages 1195–1201, 2001.
- [119] Patrick R. Jordan and Michael P. Wellman. Trading agent competition: Ad auctions, 2009. <http://www.w3.org/DesignIssues/Fractal.html>.
- [120] Adam I. Juda and David C. Parkes. An options-based method to solve the composability problem in sequential auctions. In *Agent-Mediated Electronic Commerce VI*, pages 44–58. Springer, 2006.
- [121] Adam I. Juda and David C. Parkes. The sequential auction problem on ebay: An empirical analysis and a solution. In *Proc. 7th ACM Conf. on Electr. Commerce*, pages 180–189. ACM Press, June 2006.
- [122] Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–292, 1979.
- [123] O. Kaser and D. Lemire. Tag-cloud drawing: Algorithms for cloud visualization. In *WWW'07 Workshop on Tagging and Metadata for Social Information Organization*, 2007.

- [124] M. Kearns and Ortiz. The penn-lehman automated trading project. *IEEE Intelligent Systems*, Nov–Dec, 2003.
- [125] Sheila Kinsella, Adriana Budura, Gleb Skobeltsyn, Sebastian Michel, John G. Breslin, and Karl Aberer. From web 1.0 to web 2.0 and back -: how did your grandma use to tag? In *WIDM '08: Proceeding of the 10th ACM workshop on Web information and data management*, pages 79–86. ACM, 2008.
- [126] M. Klein, P. Faratin, H. Sayama, and Y. Bar-Yam. Negotiating complex contracts. *Group Decision and Negotiation*, 12:111–125, 2003.
- [127] Paul Klemperer. Auction theory: A guide to the literature. *J. of Econ. Surveys*, 13(3):227–286, 1999.
- [128] S. Kraus and O. Schechter. Strategic-negotiation for sharing a resource between two agents. In *Computational Intelligence*, , 19(1):9-41, 2003.
- [129] Sarit Kraus. *Strategic Negotiation in Multiagent Environments*. MIT Press, Cambridge, Massachusetts, 2001.
- [130] Vijay Krishna. *Auction Theory*. Academic Press, 2002.
- [131] Ravi Kumar, Jasmine Novak, and Andrew Tomkins. Structure and evolution of online social networks. In *KDD '06: Pro. of the 12th ACM SIGKDD Int. Conf. on Knowledge Discovery and data mining*, pages 611–617. ACM, 2006.
- [132] Byron Y-L. Kuo, Tomas Hentrich, Benjamin M. Good, and Mark D. Wilkinson. Tag clouds for summarizing web search results. In *Proceedings of the 16th Int. World Wide Web Conference (WWW 2007)*, pages 1203–1204. ACM Press, 2007.
- [133] Sébastien Lahaie and David C. Parkes. Applying learning algorithms to preference elicitation. In *Proc. of 5th ACM Conference on Electronic Commerce*, pages 180–188. ACM Press, 2004.
- [134] G. Lai, C. Li, K. Sycara, and J.A. Giampapa. Literature review on multi-attribute negotiations. *Technical Report CMU-RI-TR-04-66, Robotics Institute, Carnegie Mellon University*, 2004.
- [135] Guoming Lai, Cuihong Li, and Katia Sycara. A general model for pareto optimal multi-attribute negotiations. In *Second Int. Workshop on Rational, Robust, and Secure Negotiations (RRS'06)*. Springer Computational Intelligence, 2008.
- [136] Guoming Lai, Katia Sycara, and Cuihong Li. A decentralized model for multi-attribute negotiations. In *ICEC '06: Proceedings of the 8th international conference on Electronic commerce*, pages 3–10, New York, NY, USA, 2006. ACM.
- [137] Guoming Lai, Katia Sycara, and Cuihong Li. A decentralized model for automated multi-attribute negotiations with incomplete information and general utility functions. *Multiagent Grid Systems*, 4(1):45–65, 2008.

- [138] Guoming Lai, Katia Sycara, and Cuihong Li. A decentralized model for multi-attribute negotiations with incomplete information and general utility functions. In *Proc. of Second Int. Workshop on Rational, Robust, and Secure Negotiations (RRS'06)*. Springer Series Computational Intelligence, 2008.
- [139] Jerome Lang. Logical preference representation and combinatorial voting. *Annals of Mathematics and Artificial Intelligence*, 42(1-3):37-71, 2004.
- [140] S. L. Lauritzen. *Graphical Models*. Oxford University Press, Oxford UK, 1996.
- [141] Jure Leskovec, Lada A. Adamic, and Bernardo A. Huberman. The dynamics of viral marketing. *ACM Trans. Web*, 1(1), 2007.
- [142] Kevin Leyton-Brown, Mark Pearson, and Yoav Shoham. Towards a universal test suite for combinatorial auction algorithms. In *Proc. of the ACM Conference on Electronic Commerce (ACM-EC'00)*, Minneapolis, Minnesota, 2000.
- [143] Xin Li, Lei Guo, and Yihong Eric Zhao. Tag-based social interest discovery. In *WWW '08: Proceeding of the 17th international conference on World Wide Web*, pages 675–684. ACM, 2008.
- [144] R. Lin. Bilateral multi-issue contract negotiation for task redistribution using a mediation service. In *Proc. Agent Mediated Electronic Commerce VI*, New York, USA, 2004.
- [145] Greg Linden, Brent Smith, and Jeremy York. Amazon.com recommendations: Item-to-item collaborative filtering. *IEEE Internet Computing* 7(1): 76-80 (2003), 2003.
- [146] Y. Liu, R. Goodwin, and S. Koenig. Risk-averse auction agents. In *Proc. of AAMAS'03*, pages 353–360, 2003.
- [147] A. R. Lomuscio, M. Wooldridge, and N. R. Jennings. A classification scheme for negotiation in electronic commerce. *Group Decision and Negotiation*, 12:31–56, 2003.
- [148] Xudong Luo, Nicholas R. Jennings, and Nigel R. Shadbolt. Acquiring user tradeoff strategies and preferences for negotiating agents: A default-then-adjust method. *Int. J. Hum.-Comput. Stud.*, 64(4):304–321, 2006.
- [149] Xudong Luo, Nick R. Jennings, and Nigel Shadbolt. Acquiring tradeoff preferences for automated negotiations: A case study. In *Agent-Mediated Electronic Commerce V: Designing Mechanisms and Systems*, volume 3048, pages 37–55. Springer Lecture Notes in Computer Science, 2004.
- [150] Azarakhsh Malekian, Chi-Chao Chang, Ravi Kumar, and Grant Wang. Optimizing query rewrites for keyword-based advertising. In *EC '08: Proc. of the 9th ACM conference on Electronic commerce*, pages 10–19. ACM Press, 2008.
- [151] C.D. Manning and H. Schutze. *Foundations of statistical natural language processing*. MIT Press, London, 2002.

- [152] Cameron Marlow, Mor Naaman, Danah Boyd, and Marc Davis. Position paper, tagging, taxonomy, flickr, article, toread. In *Collaborative Web Tagging Workshop at WWW'06, Edinburgh, UK*, 2006.
- [153] A. Mas-Collel, M.D. Whinston, and J.R. Green. *Microeconomic Theory*. Oxford University Press, 1995.
- [154] Adam Mathes. Folksonomies: Cooperative classification and communication through shared metadata, 2004. <http://www.adammathes.com/academic/computer-mediated-communication/folksonomies.html>.
- [155] E. Michlmayr and S. Cayzer. Learning user profiles from tagging data and leveraging them for personalized information access. In *WWW'07 Workshop on Tagging and Metadata for Social Information Organization*, 2007.
- [156] Peter Mika. Ontologies are us: A unified model of social networks and semantics. In *Proc. of the 4th Int. Semantic Web Conference (ISWC'05)*. Springer LNCS vol. 3729, 2005.
- [157] Alexander Mikroyannidis. Towards a social semantic web. *IEEE Computer Magazine*, Nov.'07: 113–115, 2007.
- [158] P. Milgrom. *Putting Auction Theory to Work*. Cambridge University Press, 2004.
- [159] P. Milgrom and R.J. Weber. A theory of auctions and competitive bidding - part 2. *Economic Theory of Auctions*, 2000.
- [160] Melanie Mitchell. *Complexity: A guided tour*. Oxford University Press, 2009.
- [161] Lonneke Mous, Valentin Robu, and Han La Poutr . Can priced options solve the exposure problem in sequential auctions? In *ACM SIGEcom Exchanges* 7(2). ACM Press, 2008.
- [162] Lonneke Mous, Valentin Robu, and Han La Poutr . Using priced options to solve the exposure problem in sequential auctions. In *Proc. of Agent-Mediated Electronic Commerce (AMEC'08)*. Springer LNCS (to appear), 2008.
- [163] R. M. Coehoorn N. R. Jennings. Learning an opponent's preferences to make effective multi-issue negotiation tradeoffs. In *Proc. 6th Int Conf. on E-Commerce, Delft*, 2004.
- [164] M. E. J. Newman. Fast algorithm for detecting community structure in networks. *Phys. Rev.*, E 69, 066133, 2004.
- [165] M. E. J. Newman and M. Girvan. Finding and evaluating community structure in networks. *Phys. Rev.*, E 69, 2004.
- [166] M.E.J. Newman. Power laws, pareto distributions and zipf's law. *Contemporary Physics*, 46:323–351, 2005.

- [167] T.D. Nguyen and N.R. Jennings. Coordinating multiple concurrent negotiations. In *3rd Int. Conf. on Aut. Agents & Multi-Agent Systems (AAMAS)*, New York, pages 1064–1071, 2004.
- [168] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [169] Noam Nisan. Bidding languages for combinatorial auctions. In *Cramton et al., Combinatorial Auctions*. MIT Press, 2006.
- [170] Steffen Oldenburg, Martin Garbe, and Clemens Cap. Similarity cross-analysis of tag / co-tag spaces in social classification systems. In *SSM '08: Proceeding of the 2008 ACM workshop on Search in social media*, pages 11–18. ACM, 2008.
- [171] H.J. Paarsch and H. Hong. *An Introduction to the Structural Econometrics of Auction Data*. MIT Press, 2006.
- [172] J. Pearl. *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan Kaufmann, San Mateo, California, 1988.
- [173] Judea Pearl. *Causality: Models, Reasoning and Inference*. Cambridge University Press, Cambridge, 2000.
- [174] Chris Preist, Claudio Bartolini, and Ivan Phillips. Algorithm design for agents which participate in multiple simultaneous auctions. *Lecture Notes in Computer Science*, 2003: 139–154, 2001.
- [175] P. Maes R. Gutman. Agent-mediated integrative negotiation for retail electronic commerce. In *Agent Mediated Electronic Commerce, Springer LNAI vol. 1571*, pages 70–90, 1998.
- [176] Matthew Rabin and Richard H. Thaler. Risk aversion. *Journal of economic perspectives*, 15(1):219–232, 2001.
- [177] Emilee Rader and Rick Wash. Influences on tag choices in del.icio.us. In *CSCW '08: Proceedings of the ACM 2008 conference on Computer supported cooperative work*, pages 239–248. ACM, 2008.
- [178] I. Rahwan, S. D. Ramchurn, N. R. Jennings, McBurney P., S. Parsons, , and Sonenberg L. Argumentation-based negotiation. *Knowledge Engineering Review*, 2004.
- [179] H. Raiffa. *The art and science of negotiation*. Harvard University Press, Cambridge, Massachussets USA, 1982.
- [180] Howard Raiffa. *Lectures on negotiation analysis*. PON Books, Harvard Law School, 1996.
- [181] S.D. Ramchurn, C. Sierra, L. Godo, and N.R. Jennings. Negotiating using rewards. In *Proc. of the 5th Int. Joint Conf. on Autonomous Agents and Multiagent Systems (AAMAS)*. ACM Press, 2006.

- [182] Tye Rattenbury, Nathaniel Good, and Mor Naaman. Towards automatic extraction of event and place semantics from flickr tags. In ACM Press, editor, *Proceedings of SIGIR'07, Amsterdam, The Netherlands*, pages 103–110, 2007.
- [183] Jonathan Wilkenfeld Raz Lin, Sarit Kraus and James Barry. An automated agent for bilateral negotiation with bounded rational agents with incomplete information. In *Proc. of the ECAI*, 2006.
- [184] D.M. Reeves, M.P. Wellman, J.K. MacKie-Mason, and A. Osepayshvili. Exploring bidding strategies for market based scheduling. *Decision Support Systems*, 39:67–85, 2005.
- [185] V. Robu and J. A. La Poutré. Retrieving the structure of utility graphs used in multi-item negotiation through collaborative filtering of aggregate buyer preferences. In *Proc. of the 2nd International Workshop on Rational, Robust and Secure Negotiations in Multi-Agent Systems (RRS 2006), Hakodate, Japan*. Springer Lecture Notes in Computational Intelligence, 2008.
- [186] V. Robu, D.J.A. Somefun, and J. A. La Poutré. Modeling complex multi-issue negotiations using utility graphs. In *Proc. of the 4th Int. Conf. on Autonomous Agents & Multi Agent Systems (AAMAS), Utrecht, The Netherlands, 2005*, pages 280–287. ACM Press, 2005.
- [187] Valentin Robu and Han La Poutré. Designing bidding strategies in sequential auctions for risk averse agents. In *Agent-Mediated Electronic Commerce (AMEC'07)*, volume 13, pages 76–89. Springer LNBIP, 2007.
- [188] Alex Rogers, Esther David, Terry R. Payne, and Nicholas R. Jennings. An advanced bidding agent for advertisement selection on public displays. In *AAMAS '07: Proc. of the 6th Int. Joint Conf. on Aut. Agents and Multiagent Systems*, pages 1–8. ACM, 2007.
- [189] J.S. Rosenschein and G. Zlotkin. *Rules of Encounter*. MIT Press, 1994.
- [190] Michael H. Rothkopf, Aleksandar Pekec, and Ronald M. Harstad. Computationally manageable combinatorial auctions. *Management Science*, 44(8):1131–1147, 1998.
- [191] A. Rucinski S. Janson, T. Luczak. *Random Graphs*. Wiley, New York, USA, 2000.
- [192] N. Sadeh, J. Eriksson, N. Finne, and S. Janson. Tac-03: A supply-chain trading competition. *AI Magazine*, 24(1):92–94, 2003.
- [193] T. Sandholm and V. Lesser. Leveled commitment contracts and strategic breach. *Games and Economic Behavior*, 35(1-2):212–270, April 2001.
- [194] Tuomas Sandholm. Algorithm for optimal winner determination in combinatorial auctions. *Artificial Intelligence*, 135(1-2):1–54, 2002.

- [195] Tuomas Sandholm and Victor Lesser. Leveled-commitment contracting: a backtracking instrument for multiagent systems. *AI Magazine*, 23(3):89–100, Fall 2002.
- [196] Kristian Hammond Sanjay Sood, Sara Owsley and Larry Birnbaum. Tagassist: Automatic tag suggestion for blog posts. In Nicolas Nicolov, Natalie Glance, Eytan Adar, Matthew Hurst, Mark Liberman, James H. Martin, and Franco Salvetti, editors, *Proceedings of the 1st International Conference on Weblogs and Social Media*, Boulder, CO, U.S.A., March 26-28 2007. <http://www.icwsm.org>.
- [197] Badrul Sarwar, George Karypis, Joseph Konstan, and John Riedl. Item-based collaborative filtering recommendation algorithms. In *Tenth International WWW Conference (WWW10)*, Hong Kong, 2001.
- [198] Ralf Schenkel, Tom Crecelius, Mouna Kacimi, Sebastian Michel, Thomas Neumann, Josiane X. Parreira, and Gerhard Weikum. Efficient top-k querying over social-tagging networks. In *SIGIR '08: Proceedings of the 31st annual international ACM SIGIR conference on Research and development in information retrieval*, pages 523–530. ACM, 2008.
- [199] C. Schmitz, M. Grahl, A. Hotho, G. Stumme, C. Cattuto, A. Baldassarri, V. Loreto, and V.D.P. Servidio. Network properties of folksonomies. In *WWW'07 Workshop on Tagging and Metadata for Social Information Organization*, 2007.
- [200] Shilad Sen, Shyong K. Lam, Al Mamunur Rashid, Dan Cosley, Dan Frankowski, Jeremy Osterhouse, F. Maxwell Harper, and John Riedl. tagging, communities, vocabulary, evolution. In *CSCW '06: Proc. of the 2006 20th Conference on Computer supported cooperative work*, pages 181–190. ACM Press, 2006.
- [201] E. Shakshuki and S. Abu-Draz. Agent-mediated e-commerce system. In *International Conference on Advanced Networking and Applications (AINA'05)*, vol. 2, pages 739–744. IEEE Computer Society, 2005.
- [202] Kaikai Shen and Lide Wu. Folksonomy as a complex network, 2005. <http://arxiv.org/abs/cs.IR/0509072>.
- [203] Ajita John Shreeharsh Kelkar and Doree Seligmann. An activity-based perspective of collaborative tagging. In Nicolas Nicolov, Natalie Glance, Eytan Adar, Matthew Hurst, Mark Liberman, James H. Martin, and Franco Salvetti, editors, *Proceedings of the 1st International Conference on Weblogs and Social Media*, Boulder, CO, U.S.A., 2007.
- [204] Carles Sierra, Peyman Faratin, and Nicholas R. Jennings. Deliberative automated negotiators using fuzzy similarities. In *EUSFLAT-ESTYLF Joint Conf.*, pages 155–158, 1999.
- [205] P. Skobelev, A. Glaschenko, I. Grachev, and S. Inozemstev. Magenta technology: case studies of magenta i-scheduler for road transportation. In *Proceedings of AA-MAS'06 (Industry Track)*, Hakodate, Japan, pages 1385–1392. ACM Press, 2006.

- [206] R. V. Sole. Syntax for free? *Nature*, 434:289, 2005.
- [207] D.J.A. Somefun, T.B. Klos, and J.A. La Poutré. Online learning of aggregate knowledge about nonlinear preferences applied to negotiating prices and bundles. In *Proc. 6th Int Conf. on E-Commerce, Delft*, pages 361–370, 2004.
- [208] K. Somefun, E. Gerding, S. Bohte, and H. La Poutre. Automated negotiation and bundling of information goods. In *Agent-Mediated Electronic Commerce V*. Springer Lecture Notes in AI, 2003.
- [209] Yang Song, Lu Zhang, and C. Lee Giles. A sparse gaussian processes classification framework for fast tag suggestions. In *CIKM '08: Proceeding of the 17th ACM conference on Information and knowledge management*, pages 93–102. ACM, 2008.
- [210] Luc Steels. The evolution of communication systems by adaptive agents. In *Adaptive agents and multi-agent systems*, pages 125–140. Springer LNAI, 2004.
- [211] M. Strobel. Effects of electronic markets on negotiation processes. In *IBM Research Technical Report, Zurich*, 2000.
- [212] Shankara B Subramanya and Huan Liu. Socialtagger - collaborative tagging for blogs in the long tail. In *SSM '08: Proceeding of the 2008 ACM workshop on Search in social media*, pages 19–26. ACM, 2008.
- [213] Panagiotis Symeonidis, Alexandros Nanopoulos, and Yannis Manolopoulos. Tag recommendations based on tensor dimensionality reduction. In *RecSys '08: Proceedings of the 2008 ACM conference on Recommender systems*, pages 43–50, New York, NY, USA, 2008. ACM.
- [214] Martin N. Szomszor, Iván Cantador, and Harith Alani. Correlating user profiles from multiple folksonomies. In *HT '08: Proceedings of the nineteenth ACM conference on Hypertext and hypermedia*, pages 33–42. ACM, 2008.
- [215] R. L. Rivest T. H. Cormen, C. E. Leiserson. *Introduction to algorithms*. The MIT Press, 1989.
- [216] Pieter Jan 't Hoen, Valentin Robu, and Han La Poutré. *Decommitment in a Competitive Multi-Agent Transportation Setting*, pages 409–431. Whitestein Series in Software Agent Technologies. Birkhauser, Springer, 2005.
- [217] P.J. 't Hoen and J.A. La Poutré. Repeated auctions with complementarities. In *Proc. of 7th Agent-Mediated Electronic Commerce Workshop (AMEC'05), Utrecht*, 2005.
- [218] Leigh Tesfatsion and Kenneth Judd. *Handbook of Computational Economics, vol. 2: Agent-Based Computational Economics*. Elsevier, North Holland, 2006.
- [219] D. Koller U. Chajewska and D. Ormoneit. Learning an agent's utility function by observing behavior. In *18th International Conference on Machine Learning (ICML '01)*, pages 35–42, 2001.

- [220] J.A. La Poutré V. Robu. Retrieving the structure of utility graphs used in multi-issue negotiation through collaborative filtering. In *1st International Workshop on Computational Social Choice, Amsterdam*. ILLC, Univ. of Amsterdam, The Netherlands, 2006.
- [221] P.V. van Amstel and A.R. Goor. *Van Logistiek naar Supply Chain Management*. Kluwer Publishers, 2003.
- [222] S. van der Putten, V. Robu, J.A. La Poutré, A. Jorritsma, and M. Gal. Automating supply chain negotiations using autonomous agents: a case study in transportation logistics. In *Proc. 5th Int. Joint Conf. on Autonomous Agents and Multi Agent Systems (AAMAS'06), Industry Track*, pages 1506–1513. ACM Press, 2006.
- [223] Ioannis A. Vetsikas and Nicholas R. Jennings. Bidding strategies for realistic multi-unit sealed-bid auctions. In *Proc. of the 23rd Conference on Artificial Intelligence (AAAI'08), Chicago, USA*, 2008.
- [224] Ioannis A. Vetsikas and Nicholas R. Jennings. Towards agents participating in realistic multi-unit sealed-bid auctions. In *Proc. 7th Int. Conf. on Autonomous Agents and Multi-agent Systems (AAMAS'08), Estoril, Portugal*, pages 1621–1624, 2008.
- [225] Quoc Bao Vo and Lin Padgham. Searching for joint gains in automated negotiations based on multi-criteria decision making theory. In *AAMAS '07: Proceedings of the 6th international joint conference on Autonomous agents and multiagent systems*, pages 1–8, New York, NY, USA, 2007. ACM.
- [226] P. Vytelingum, R. K. Dash, E. David, and N. R. Jennings. A risk-based bidding strategy for continuous double auctions. In *Proc. of 16th ECAI, Valencia, Spain*, pages 79–83, 2004.
- [227] P. Vytelingum, R. K. Dash, I. A. Vetsikas, and N. R. Jennings. Whitedolphin: A tac travel agent. In *AAAI Workshop on Trading Agent Design and Analysis*, 2007.
- [228] Mitchell Waldrop. *Complexity: The Emerging Science at the Edge of Order and Chaos*. Simon and Schuster, 1992.
- [229] Duncan Watts and Steve Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393(6684):440–442, 1998.
- [230] M.P. Wellman, A. Greenwald, and P. Stone. *Autonomous Bidding Agents: Strategies and Lessons from the Trading Agent Competition*. MIT Press, 2007.
- [231] Michael Wooldridge and Nicholas R. Jennings. Intelligent agents: Theory and practice. *The Knowledge Engineering Review*, 10(2):115–152, 1995.
- [232] N. Shadbolt H. Leung J. H. Lee X. Luo, N. R. Jennings. A fuzzy constraint based model for bilateral multi-issue negotiations in semi-competitive environments. *Artificial Intelligence Journal*, 142 (1-2):53–102, 2003.

- [233] Shohei Yoshikawa, Takahiko Kamiryo, Yoshiaki Yasumura, and Kuniaki Uehara. Strategy acquisition of agents in multi-issue negotiation. In *WI '06: Proceedings of the 2006 IEEE/WIC/ACM International Conference on Web Intelligence*, pages 933–939, Washington, DC, USA, 2006. IEEE Computer Society.
- [234] Valentina Zanardi and Licia Capra. Social ranking: uncovering relevant content using tag-based recommender systems. In *RecSys '08: Proceedings of the 2008 ACM conference on Recommender systems*, pages 51–58. ACM, 2008.
- [235] A. Zollers. Emerging motivations for tagging: Expression, performance and activism. In *WWW'07 Workshop on Tagging and Metadata for Social Information Organization*, 2007.

Summary

Multi-agent systems represent an important, emerging area of research at the border between artificial intelligence and distributed systems, on one hand, and economics and game theory on the other. Briefly defined, agents are autonomous, pro-active software programs that can reason and take decisions on behalf of their human owners. One of the main application areas for multi-agent systems is the automation of electronic markets, such as electronic negotiation and auction environments. Some important application areas (that have been considered, in depth, as part of this thesis) are the automation of complex negotiations in online electronic commerce and task allocation between multiple companies in distributed logistics.

This thesis studies several important, open problems in agent-based electronic markets. The first of these is how complex preferences and utility functions can be modelled and used to design efficient strategies for bilateral negotiation and auction situations. Here, we distinguish between two classes of preferences: combinatorial preferences (over combinations of a large number of items or issues) and preferences towards risk (i.e. how risk averse agent is when taking decisions in an uncertain environment). The second important aspect we study is the strategic reasoning, especially reasoning of agents participating in a sequence of auctions with complementarities. For this case, we propose a novel priced options mechanism, that can reduce the exposure problem facing bidders participating in such auctions. Finally, the third important problem studied in this thesis is how collaboration can emerge in a system composed of many autonomous, self-interested agents. For this part, we use large scale empirical data from two social web applications: collaborative tagging and sponsored search markets.

The contributions to the literature, per each chapter, can be summarised as follows. Chapter 2 considers the problem of modeling bilateral, multi-attribute negotiations in environments with incomplete preference information, but in which the preferences of the negotiating agents can be represented as linearly additive utility functions. We propose an algorithm that allows agents to use incomplete preference information in automatic negotiation and reach jointly profitable agreements.

Chapter 3 also considers the problem of automated multi-issue or multi-item negotiation, but for the case that there are utility interdependencies between items sold, a problem which is known to be considerably more challenging than the linearly independent issues case, and

for which very few computationally efficient models were known to exist. We propose a novel utility graphs formalism and show how it can be used to efficiently encode and learn opponent utilities in complex, multi-item negotiations. In an extension of this work, we show that the initial, starting structure of such utility graphs can be approximated by using collaborative filtering on concluded negotiation data. This provides an important link between collaborative filtering techniques used in electronic commerce and multi-item negotiation.

Another important line of work for the thesis (considered in Chapter 4) is designing efficient bidding strategies in sequential auction settings, for risk-averse agents. While the problem of designing efficient auction bidding strategies in uncertain, sequential environments was well known, previous literature did not consider how risk aversion affects an agent's optimal bidding policy. The motivation for this problem came from the real business case regarding transportation logistics: in real life, transportation providers are reluctant to use bidding strategies that could lead to large losses, even if they maximize their expected utility.

Chapter 5 considers the exposure problem that agents with complementary valuations over combinations of goods face when bidding in sequential auctions in which these goods are sold independently. In order to solve this problem, we study a more complex market mechanism: selling priced options for the goods, instead of the goods themselves. Our work builds on and extends the concept of non-priced options proposed by Juda & Parkes, and it shows how options could be priced in order to bring the maximum benefit to both buyers and sellers participating in such an uncertain market environment.

The third part of the thesis considers a somewhat different issue: how collaboration and social preference form in large systems composed of many self-interested agents. In order to do this, we studied two such systems: a large dataset of sponsored search data, provided from a project with Microsoft Research, and collaborative tagging dataset (obtained from Del.icio.us/ Yahoo). Chapter 6 looks at the issue of how stable vocabularies form, in the absence of a central controller, in large-scale, collaborative tagging systems. This chapter uses techniques first developed in the field of complex systems theory. Chapter 7 uses similar complex systems techniques, but now applied to sponsored search markets. This chapter also provides a link to the auction-based approach, used in the second part of the thesis.

The dissertation is concluded by an appendix describing an industrial case study, which investigated the applicability of some of the techniques described in the more theoretical chapters (especially the auction-based techniques in Chapters 4 and 5) to distributed transportation logistics. This work was conducted in collaboration with Vos Logistics Organizing, Nijmegen (VLO). Based on a real-world scenario provided by VLO, we built a multi-agent platform in which transportation orders can be allocated dynamically between different companies, through a system of dynamic, distributed auctions.

Samenvatting

Multi-agent systemen is een belangrijk opkomend terrein van onderzoek, op het grensvlak van kunstmatige intelligentie, gedistribueerde systemen, de economie en de speltheorie. Kort gedefiniëerd zijn agenten autonome, pro-actieve software programmas die zelfstandig kunnen redeneren, en vervolgens beslissingen kunnen nemen namens hun menselijke eigenaren. Een van de voornaamste toepassingsgebieden voor multi-agent systems is het automatiseren van elektronische markten, zoals elektronische onderhandelingen en veiling-omgevingen. Een belangrijk toepassingsgebied, zoals diepgaand wordt bestudeerd in dit proefschrift, is de automatisering van complexe onderhandelingen in online e-commerce en taakverdeling tussen meerdere bedrijven in gedistribueerde logistiek.

In dit proefschrift zijn een aantal belangrijke en open problemen in agent-gebaseerde elektronische markten bestudeerd. Ten eerste is bestudeerd hoe complexe preferenties en utility-functies kunnen worden gemodelleerd en gebruikt om efficiënte strategieën te ontwikkelen voor bilaterale onderhandelingen en veilingen. We onderscheiden hier twee klassen van preferenties: combinatorische preferenties (over combinaties van een groot aantal objecten of afwegingen) en risico-preferenties (bijvoorbeeld hoe risico-mijdend een agent is bij het nemen van beslissingen in een onzekere omgeving). Een tweede belangrijk aspect dat we bestuderen, is dat van strategisch beredeneren, en dan met name wanneer agenten deelnemen aan een serie van veilingen met complementariteiten. Voor deze situatie introduceren we een nieuw mechanisme gebaseerd op het prijzen van opties. Dit mechanisme reduceert het exposure probleem voor deelnemers aan dergelijke sequentiële veilingen. Het derde bestudeerde probleem, ten slotte, is hoe samenwerking kan ontstaan in een systeem dat bestaat uit vele autonome en belanghebbende agenten. We bestuderen hier twee grote empirische datasets, afkomstig van twee sociale web applicaties: collaboratieve “tagging” en de markt van sponsored search.

De bijdragen aan de literatuur, per hoofdstuk, kunnen als volgt worden samengevat. Hoofdstuk 2 beschouwt het probleem van het modelleren van bilaterale onderhandelingen over meerdere eigenschappen, in omgevingen met incomplete informatie, maar waarin de preferenties van de onderhandelende agenten kunnen worden gerepresenteerd als additieve utility functies. We introduceren een algoritme dat het agenten mogelijk maakt om incomplete informatie te gebruiken in automatische onderhandelingen en zo voor allen profitabele overeenkomsten te bereiken.

Hoofdstuk 3 beschouwt ook het probleem van automatische onderhandelingen over meerdere eigenschappen of objecten, maar voor het geval wanneer er interdependencies bestaan tussen de objecten. Dit probleem staat bekend als aanzienlijk complexer dan het lineair onafhankelijke geval, en er bestaan zeer weinig computationeel efficiënte modellen. We introduceren een nieuw formalisme voor utiliteits grafen, en we laten zien hoe dit kan worden gebruikt om de utilities van een opponenent efficient te leren en te coderen in complexe onderhandelingen over meerdere objecten. In een uitbreiding op dit werk laten we zien dat de initiële toestandsstructuur van dergelijke utiliteits grafen kan worden benaderd door gebruik te maken van collaboratief filteren op de gegevens van voorgaande onderhandelingen. Deze uitbreiding slaat tevens een brug tussen enerzijds collaboratief filteren in e-commerce toepassingen, en anderzijds het onderhandelen over meerdere objecten.

Een andere belangrijke lijn in dit proefschrift, zoals beschouwd in Hoofdstuk 4, is het ontwerp van efficiënte biedstrategieën in sequentiële veilingen, specifiek voor risico-mijdende agenten. Hoewel het probleem van het ontwerpen van efficiënte biedstrategieën in onzekere en sequentiële ongevingen bekend was in de literatuur, was niet bekend hoe risico-mijdende preferenties van invloed zijn op de optimale biedstrategie van een agent. De motivatie voor dit probleem kwam uit de transport-logistiek business case: in de praktijk zijn transporteurs terughoudend in het gebruik van biedstrategieën die mogelijk tot grote verliezen zouden kunnen leiden, zelfs als deze strategie hun gemiddelde, verwachte utility maximaliseert.

In Hoofdstuk 5 beschouwen we het exposure probleem van agenten die complementaire waarderingen hebben over combinaties van goederen, wanneer deze agenten de goederen kunnen verkrijgen in sequentiële, onafhankelijke veilingen. Om dit probleem op te lossen, bestuderen we een meer complex markt-mechanisme: het verkopen van opties op de goederen, in plaats van de goederen zelf. Dit werk bouwt voort op, en breidt uit, het concept van niet-geprijsde opties van Juda & Parkes, en het laat zien hoe opties geprijsd moeten worden om de waarde te maximaliseren voor zowel kopers als verkopers in een dergelijke onzekere marktomgeving.

Het derde deel van dit proefschrift beschouwt een ander aspect van multi-agent systemen: hoe samenwerking en social preferentie tot stand komen in grote systemen die bestaan uit zelfstandige agenten die uitsluitend handelen in hun eigen belang. Hiertoe bestuderen we twee van dergelijke systemen: een grote dataset van sponsored search data, verkregen in het kader van een project met Microsoft Research, en een “collaborative tagging” dataset, verkregen van Del.icio.us/Yahoo. Hoofdstuk 6 bestudeert de vorming van stabiele vocabulaires, in de afwezigheid van centrale sturing, in een groot collaborative tagging systeem. Dit hoofdstuk gebruikt technieken oorspronkelijk ontwikkeld in het veld van de complexe systeem theorie. Hoofdstuk 7 gebruikt soortgelijke technieken, maar nu toegepast op sponsored search markten. Dit hoofdstuk vormt zo ook een brug naar de veiling-gebaseerde benaderingen, zoals gebruikt in het tweede deel van dit proefschrift.

Dit proefschrift wordt besloten met een appendix waarin een industriële case studie is beschreven. Hierin wordt de toepasbaarheid bestudeerd van enkele van de technieken zoals beschreven in de meer theoretische hoofdstukken (en dan met name de veilingstechnieken van Hoofdstuk 4 en 5), in de context van gedistribueerde transportlogistiek. Dit werk is uitgevo-

erd in samenwerking met Vos Logistics Organizing, Nijmegen (VLO). Op basis van een door VLO aangereikt real-world scenario hebben we een multi-agent platform gebouwd waarin transport order dynamisch kunnen worden gealloceerd tussen verschillende bedrijven, door middel van een systeem van dynamische en gedistribueerde veilingen.

Curriculum Vitae

Valentin Robu was born in Cluj-Napoca, Transylvania, Romania in 1978. He graduated the “Emil Racoviță” High School from the same city in 1998. From 1998, he studied computer science at the Technical University of Cluj-Napoca Romania. In parallel, he also followed courses in Economics (with specialization in capital markets) at the Babeş-Bolyai University in Cluj.

In April 2002, he received a scholarship to complete a master degree in Artificial Intelligence (Dutch “doctorandus”) at the Free University of Amsterdam, the Netherlands. He obtained his master degree in August 2003, with a thesis on modeling automated negotiations in incomplete information environments (completed under the supervision of prof. dr. Catholijn Jonker and prof. dr. Jan Treur). During this time, he also worked part-time (two days/week) at Almende B.V., a software consultancy and research company in Rotterdam.

From October 2003, he was employed as a PhD student at CWI, the National Center for Mathematics and Computer Science in the Netherlands, Amsterdam, under the supervision of prof. Han La Poutré. His PhD work was carried out in the framework of the Senter-Novem project DEAL (Distributed Engine for Advanced Logistics). The overall topic of his PhD research concerns modeling preferences and strategies in agent-mediated electronic market settings, such as bilateral negotiations and distributed auctions. It resulted in several publications in international conferences, book chapters and journals (a list of which is available at the end of the introduction of this thesis). During his PhD, he was awarded travel grants to attend several international conferences (such as AAMAS '04-AAMAS'08 and WWW '07), 3 Dagstuhl seminars, as well as the Complex Systems Summer School organised by the Santa Fe Institute in 2006.

From April 2008 to March 2009, he was employed at CWI as a research fellow in the DIACoDeM project (Distributed, Intelligent Adaptive Collective Decision Making). During this time, he also worked as 0.2 FTE as a guest lecturer in the Algorithms group at the Technical University of Delft, and also supervised two master thesis projects. From March 2009, he is employed as a research fellow in the Intelligence, Agents, Multimedia Group at the University of Southampton, United Kingdom, under the supervision of prof. Nick Jennings.

List of SIKS Dissertations - SIKS Dissertatiereeks

====
1998
====

- 1998-1 Johan van den Akker (CWI)
DEGAS - An Active, Temporal Database of Autonomous Objects
- 1998-2 Floris Wiesman (UM)
Information Retrieval by Graphically Browsing Meta-Information
- 1998-3 Ans Steuten (TUD)
A Contribution to the Linguistic Analysis of Business Conversations
within the Language/Action Perspective
- 1998-4 Dennis Breuker (UM)
Memory versus Search in Games
- 1998-5 E.W.Oskamp (RUL)
Computerondersteuning bij Straftoemeting

====
1999
====

- 1999-1 Mark Sloof (VU)
Physiology of Quality Change Modelling;
Automated modelling of Quality Change of Agricultural Products
- 1999-2 Rob Potharst (EUR)
Classification using decision trees and neural nets
- 1999-3 Don Beal (UM)
The Nature of Minimax Search
- 1999-4 Jacques Penders (UM)
The practical Art of Moving Physical Objects
- 1999-5 Aldo de Moor (KUB)
Empowering Communities: A Method for the Legitimate User-Driven
Specification of Network Information Systems
- 1999-6 Niek J.E. Wijngaards (VU)
Re-design of compositional systems
- 1999-7 David Spelt (UT)
Verification support for object database design
- 1999-8 Jacques H.J. Lenting (UM)
Informed Gambling: Conception and Analysis of a Multi-Agent
Mechanism for Discrete Reallocation.

====
2000
====

- 2000-1 Frank Niessink (VU)
Perspectives on Improving Software Maintenance

- 2000-2 Koen Holtman (TUE)
Prototyping of CMS Storage Management
- 2000-3 Carolien M.T. Metselaar (UVA)
Sociaal-organisatorische gevolgen van kennistechnologie;
een procesbenadering en actorperspectief.
- 2000-4 Geert de Haan (VU)
ETAG, A Formal Model of Competence Knowledge for User Interface Design
- 2000-5 Ruud van der Pol (UM)
Knowledge-based Query Formulation in Information Retrieval.
- 2000-6 Rogier van Eijk (UU)
Programming Languages for Agent Communication
- 2000-7 Niels Peek (UU)
Decision-theoretic Planning of Clinical Patient Management
- 2000-8 Veerle Coup (EUR)
Sensitivity Analysis of Decision-Theoretic Networks
- 2000-9 Florian Waas (CWI)
Principles of Probabilistic Query Optimization
- 2000-10 Niels Nes (CWI)
Image Database Management System Design Considerations,
Algorithms and Architecture
- 2000-11 Jonas Karlsson (CWI)
Scalable Distributed Data Structures for Database Management

====
2001
====

- 2001-1 Silja Renooij (UU)
Qualitative Approaches to Quantifying Probabilistic Networks
- 2001-2 Koen Hindriks (UU)
Agent Programming Languages: Programming with Mental Models
- 2001-3 Maarten van Someren (UvA)
Learning as problem solving
- 2001-4 Evgueni Smirnov (UM)
Conjunctive and Disjunctive Version Spaces with
Instance-Based Boundary Sets
- 2001-5 Jacco van Ossenbruggen (VU)
Processing Structured Hypermedia: A Matter of Style
- 2001-6 Martijn van Welie (VU)
Task-based User Interface Design
- 2001-7 Bastiaan Schonhage (VU)
Diva: Architectural Perspectives on Information Visualization
- 2001-8 Pascal van Eck (VU)
A Compositional Semantic Structure for Multi-Agent Systems Dynamics.
- 2001-9 Pieter Jan 't Hoen (RUL)
Towards Distributed Development of Large Object-Oriented Models,

Views of Packages as Classes

2001-10 Maarten Sierhuis (UvA)
Modeling and Simulating Work Practice
BRAHMS: a multiagent modeling and simulation language
for work practice analysis and design

2001-11 Tom M. van Engers (VUA)
Knowledge Management:
The Role of Mental Models in Business Systems Design

====
2002
====

2002-01 Nico Lassing (VU)
Architecture-Level Modifiability Analysis

2002-02 Roelof van Zwol (UT)
Modelling and searching web-based document collections

2002-03 Henk Ernst Blok (UT)
Database Optimization Aspects for Information Retrieval

2002-04 Juan Roberto Castelo Valdueza (UU)
The Discrete Acyclic Digraph Markov Model in Data Mining

2002-05 Radu Serban (VU)
The Private Cyberspace Modeling Electronic Environments
inhabited by Privacy-concerned Agents

2002-06 Laurens Mommers (UL)
Applied legal epistemology;
Building a knowledge-based ontology of the legal domain

2002-07 Peter Boncz (CWI)
Monet: A Next-Generation DBMS Kernel For Query-Intensive Applications

2002-08 Jaap Gordijn (VU)
Value Based Requirements Engineering: Exploring Innovative
E-Commerce Ideas

2002-09 Willem-Jan van den Heuvel(KUB)
Integrating Modern Business Applications with Objectified Legacy Systems

2002-10 Brian Sheppard (UM)
Towards Perfect Play of Scrabble

2002-11 Wouter C.A. Wijngaards (VU)
Agent Based Modelling of Dynamics: Biological and Organisational Applications

2002-12 Albrecht Schmidt (Uva)
Processing XML in Database Systems

2002-13 Hongjing Wu (TUE)
A Reference Architecture for Adaptive Hypermedia Applications

2002-14 Wieke de Vries (UU)
Agent Interaction: Abstract Approaches to Modelling, Programming and
Verifying Multi-Agent Systems

2002-15 Rik Eshuis (UT)
Semantics and Verification of UML Activity Diagrams for Workflow Modelling

- 2002-16 Pieter van Langen (VU)
The Anatomy of Design: Foundations, Models and Applications
- 2002-17 Stefan Manegold (UVA)
Understanding, Modeling, and Improving Main-Memory Database Performance

====
2003
====

- 2003-01 Heiner Stuckenschmidt (VU)
Ontology-Based Information Sharing in Weakly Structured Environments
- 2003-02 Jan Broersen (VU)
Modal Action Logics for Reasoning About Reactive Systems
- 2003-03 Martijn Schuemie (TUD)
Human-Computer Interaction and Presence in Virtual Reality Exposure Therapy
- 2003-04 Milan Petkovic (UT)
Content-Based Video Retrieval Supported by Database Technology
- 2003-05 Jos Lehmann (UVA)
Causation in Artificial Intelligence and Law - A modelling approach
- 2003-06 Boris van Schooten (UT)
Development and specification of virtual environments
- 2003-07 Machiel Jansen (UvA)
Formal Explorations of Knowledge Intensive Tasks
- 2003-08 Yongping Ran (UM)
Repair Based Scheduling
- 2003-09 Rens Kortmann (UM)
The resolution of visually guided behaviour
- 2003-10 Andreas Lincke (UvT)
Electronic Business Negotiation: Some experimental studies on the interaction between medium, innovation context and culture
- 2003-11 Simon Keizer (UT)
Reasoning under Uncertainty in Natural Language Dialogue using Bayesian Networks
- 2003-12 Roeland Ordelman (UT)
Dutch speech recognition in multimedia information retrieval
- 2003-13 Jeroen Donkers (UM)
Nosce Hostem - Searching with Opponent Models
- 2003-14 Stijn Hoppenbrouwers (KUN)
Freezing Language: Conceptualisation Processes across ICT-Supported Organisations
- 2003-15 Mathijs de Weerd (TUD)
Plan Merging in Multi-Agent Systems
- 2003-16 Menzo Windhouwer (CWI)
Feature Grammar Systems - Incremental Maintenance of Indexes to Digital Media Warehouses
- 2003-17 David Jansen (UT)
Extensions of Statecharts with Probability, Time, and Stochastic Timing

2003-18 Levente Kocsis (UM)
Learning Search Decisions

====
2004
====

2004-01 Virginia Dignum (UU)
A Model for Organizational Interaction: Based on Agents, Founded in Logic

2004-02 Lai Xu (UvT)
Monitoring Multi-party Contracts for E-business

2004-03 Perry Groot (VU)
A Theoretical and Empirical Analysis of Approximation in Symbolic Problem Solving

2004-04 Chris van Aart (UVA)
Organizational Principles for Multi-Agent Architectures

2004-05 Viara Popova (EUR)
Knowledge discovery and monotonicity

2004-06 Bart-Jan Hommes (TUD)
The Evaluation of Business Process Modeling Techniques

2004-07 Elise Boltjes (UM)
Voorbeeldig onderwijs; voorbeeldgestuurd onderwijs, een opstap naar
abstract denken, vooral voor meisjes

2004-08 Joop Verbeek(UM)
Politie en de Nieuwe Internationale Informatiemarkt, Grensregionale
politiele gegevensuitwisseling en digitale expertise

2004-09 Martin Caminada (VU)
For the Sake of the Argument; explorations into argument-based reasoning

2004-10 Suzanne Kabel (UVA)
Knowledge-rich indexing of learning-objects

2004-11 Michel Klein (VU)
Change Management for Distributed Ontologies

2004-12 The Duy Bui (UT)
Creating emotions and facial expressions for embodied agents

2004-13 Wojciech Jamroga (UT)
Using Multiple Models of Reality: On Agents who Know how to Play

2004-14 Paul Harrenstein (UU)
Logic in Conflict. Logical Explorations in Strategic Equilibrium

2004-15 Arno Knobbe (UU)
Multi-Relational Data Mining

2004-16 Federico Divina (VU)
Hybrid Genetic Relational Search for Inductive Learning

2004-17 Mark Winands (UM)
Informed Search in Complex Games

2004-18 Vania Bessa Machado (UvA)
Supporting the Construction of Qualitative Knowledge Models

- 2004-19 Thijs Westerveld (UT)
Using generative probabilistic models for multimedia retrieval
- 2004-20 Madelon Evers (Nyenrode)
Learning from Design: facilitating multidisciplinary design teams

====
2005
====

- 2005-01 Floor Verdenius (UVA)
Methodological Aspects of Designing Induction-Based Applications
- 2005-02 Erik van der Werf (UM)
AI techniques for the game of Go
- 2005-03 Franc Grootjen (RUN)
A Pragmatic Approach to the Conceptualisation of Language
- 2005-04 Nirvana Meratnia (UT)
Towards Database Support for Moving Object data
- 2005-05 Gabriel Infante-Lopez (UVA)
Two-Level Probabilistic Grammars for Natural Language Parsing
- 2005-06 Pieter Spronck (UM)
Adaptive Game AI
- 2005-07 Flavius Frasincar (TUE)
Hypermedia Presentation Generation for Semantic Web Information Systems
- 2005-08 Richard Vdovjak (TUE)
A Model-driven Approach for Building Distributed Ontology-based Web Applications
- 2005-09 Jeen Broekstra (VU)
Storage, Querying and Inferencing for Semantic Web Languages
- 2005-10 Anders Bouwer (UVA)
Explaining Behaviour: Using Qualitative Simulation in Interactive Learning Environments
- 2005-11 Elth Ogston (VU)
Agent Based Matchmaking and Clustering - A Decentralized Approach to Search
- 2005-12 Csaba Boer (EUR)
Distributed Simulation in Industry
- 2005-13 Fred Hamburg (UL)
Een Computermodel voor het Ondersteunen van Euthanasiebeslissingen
- 2005-14 Borys Omelayenko (VU)
Web-Service configuration on the Semantic Web; Exploring how semantics meets pragmatics
- 2005-15 Tibor Bosse (VU)
Analysis of the Dynamics of Cognitive Processes
- 2005-16 Joris Graaumans (UU)
Usability of XML Query Languages
- 2005-17 Boris Shishkov (TUD)
Software Specification Based on Re-usable Business Components

- 2005-18 Danielle Sent (UU)
Test-selection strategies for probabilistic networks
- 2005-19 Michel van Dartel (UM)
Situated Representation
- 2005-20 Cristina Coteanu (UL)
Cyber Consumer Law, State of the Art and Perspectives
- 2005-21 Wijnand Derks (UT)
Improving Concurrency and Recovery in Database Systems by
Exploiting Application Semantics

====
2006
====

- 2006-01 Samuil Angelov (TUE)
Foundations of B2B Electronic Contracting
- 2006-02 Cristina Chisalita (VU)
Contextual issues in the design and use of information technology in organizations
- 2006-03 Noor Christoph (UVA)
The role of metacognitive skills in learning to solve problems
- 2006-04 Marta Sabou (VU)
Building Web Service Ontologies
- 2006-05 Cees Pierik (UU)
Validation Techniques for Object-Oriented Proof Outlines
- 2006-06 Ziv Baida (VU)
Software-aided Service Bundling - Intelligent Methods & Tools
for Graphical Service Modeling
- 2006-07 Marko Smiljanic (UT)
XML schema matching – balancing efficiency and effectiveness by means of clustering
- 2006-08 Eelco Herder (UT)
Forward, Back and Home Again - Analyzing User Behavior on the Web
- 2006-09 Mohamed Wahdan (UM)
Automatic Formulation of the Auditor's Opinion
- 2006-10 Ronny Siebes (VU)
Semantic Routing in Peer-to-Peer Systems
- 2006-11 Joeri van Ruth (UT)
Flattening Queries over Nested Data Types
- 2006-12 Bert Bongers (VU)
Interactivation - Towards an e-cology of people, our technological environment, and the arts
- 2006-13 Henk-Jan Lebbink (UU)
Dialogue and Decision Games for Information Exchanging Agents
- 2006-14 Johan Hoom (VU)
Software Requirements: Update, Upgrade, Redesign - towards a Theory of Requirements Change
- 2006-15 Rainer Malik (UU)
CONAN: Text Mining in the Biomedical Domain

- 2006-16 Carsten Riggelsen (UU)
Approximation Methods for Efficient Learning of Bayesian Networks
- 2006-17 Stacey Nagata (UU)
User Assistance for Multitasking with Interruptions on a Mobile Device
- 2006-18 Valentin Zhizhkin (UVA)
Graph transformation for Natural Language Processing
- 2006-19 Birna van Riemsdijk (UU)
Cognitive Agent Programming: A Semantic Approach
- 2006-20 Marina Velikova (UvT)
Monotone models for prediction in data mining
- 2006-21 Bas van Gils (RUN)
Aptness on the Web
- 2006-22 Paul de Vrieze (RUN)
Fundamentals of Adaptive Personalisation
- 2006-23 Ion Juvina (UU)
Development of Cognitive Model for Navigating on the Web
- 2006-24 Laura Hollink (VU)
Semantic Annotation for Retrieval of Visual Resources
- 2006-25 Madalina Drugan (UU)
Conditional log-likelihood MDL and Evolutionary MCMC
- 2006-26 Vojkan Mihajlovic (UT)
Score Region Algebra: A Flexible Framework for Structured Information Retrieval
- 2006-27 Stefano Bocconi (CWI)
Vox Populi: generating video documentaries from semantically annotated media repositories
- 2006-28 Borkur Sigurbjornsson (UVA)
Focused Information Access using XML Element Retrieval

====
2007
====

- 2007-01 Kees Leune (UvT)
Access Control and Service-Oriented Architectures
- 2007-02 Wouter Teepe (RUG)
Reconciling Information Exchange and Confidentiality: A Formal Approach
- 2007-03 Peter Mika (VU)
Social Networks and the Semantic Web
- 2007-04 Jurriaan van Diggelen (UU)
Achieving Semantic Interoperability in Multi-agent Systems: a dialogue-based approach
- 2007-05 Bart Schermer (UL)
Software Agents, Surveillance, and the Right to Privacy: a Legislative Framework for Agent-enabled Surveillance
- 2007-06 Gilad Mishne (UVA)
Applied Text Analytics for Blogs
- 2007-07 Natasa Jovanovic' (UT)
To Whom It May Concern - Addressee Identification in Face-to-Face Meetings

- 2007-08 Mark Hoogendoorn (VU)
Modeling of Change in Multi-Agent Organizations
- 2007-09 David Mobach (VU)
Agent-Based Mediated Service Negotiation
- 2007-10 Huib Aldewereld (UU)
Autonomy vs. Conformity: an Institutional Perspective on Norms and Protocols
- 2007-11 Natalia Stash (TUE)
Incorporating Cognitive/Learning Styles in a General-Purpose Adaptive Hypermedia System
- 2007-12 Marcel van Gerven (RUN)
Bayesian Networks for Clinical Decision Support: A Rational Approach to Dynamic Decision-Making under Uncertainty
- 2007-13 Rutger Rienks (UT)
Meetings in Smart Environments; Implications of Progressing Technology
- 2007-14 Niek Bergboer (UM)
Context-Based Image Analysis
- 2007-15 Joyca Lacroix (UM)
NIM: a Situated Computational Memory Model
- 2007-16 Davide Grossi (UU)
Designing Invisible Handcuffs. Formal investigations in Institutions and Organizations for Multi-agent Systems
- 2007-17 Theodore Charitos (UU)
Reasoning with Dynamic Networks in Practice
- 2007-18 Bart Orriens (UvT)
On the development and management of adaptive business collaborations
- 2007-19 David Levy (UM)
Intimate relationships with artificial partners
- 2007-20 Slinger Jansen (UU)
Customer Configuration Updating in a Software Supply Network
- 2007-21 Karianne
(UU)
Fast diffusion and broadening use: A research on residential adoption and usage of
broadband internet in the Netherlands between 2001 and 2005
- 2007-22 Zlatko Zlatev (UT)
Goal-oriented design of value and process models from patterns
- 2007-23 Peter Bama (TUE)
Specification of Application Logic in Web Information Systems
- 2007-24 Georgina Ramirez Camps (CWI)
Structural Features in XML Retrieval
- 2007-25 Joost Schalken (VU)
Empirical Investigations in Software Process Improvement

====
2008
====

- 2008-01 Katalin Boer-Sorbn (EUR)
Agent-Based Simulation of Financial Markets: A modular, continuous-time approach

- 2008-02 Alexei Sharpanskykh (VU)
On Computer-Aided Methods for Modeling and Analysis of Organizations
- 2008-03 Vera Hollink (UVA)
Optimizing hierarchical menus: a usage-based approach
- 2008-04 Ander de Keijzer (UT)
Management of Uncertain Data - towards unattended integration
- 2008-05 Bela Mutschler (UT)
Modeling and simulating causal dependencies on process-aware information systems from a cost perspective
- 2008-06 Arjen Hommersom (RUN)
On the Application of Formal Methods to Clinical Guidelines, an Artificial Intelligence Perspective
- 2008-07 Peter van Rosmalen (OU)
Supporting the tutor in the design and support of adaptive e-learning
- 2008-08 Janneke Bolt (UU)
Bayesian Networks: Aspects of Approximate Inference
- 2008-09 Christof van Nimwegen (UU)
The paradox of the guided user: assistance can be counter-effective
- 2008-10 Wouter Bosma (UT)
Discourse oriented summarization
- 2008-11 Vera Kartseva (VU)
Designing Controls for Network Organizations: A Value-Based Approach
- 2008-12 Jozsef Farkas (RUN)
A Semiotically Oriented Cognitive Model of Knowledge Representation
- 2008-13 Caterina Carraciolo (UVA)
Topic Driven Access to Scientific Handbooks
- 2008-14 Arthur van Bunningen (UT)
Context-Aware Querying; Better Answers with Less Effort
- 2008-15 Martijn van Otterlo (UT)
The Logic of Adaptive Behavior: Knowledge Representation and Algorithms for the Markov Decision Process Framework in First-Order Domains.
- 2008-16 Henriette van Vugt (VU)
Embodied agents from a user's perspective
- 2008-17 Martin Op 't Land (TUD)
Applying Architecture and Ontology to the Splitting and Allying of Enterprises
- 2008-18 Guido de Croon (UM)
Adaptive Active Vision
- 2008-19 Henning Rode (UT)
From Document to Entity Retrieval: Improving Precision and Performance of Focused Text Search
- 2008-20 Rex Arendsen (UVA)
Geen bericht, goed bericht. Een onderzoek naar de effecten van de introductie van elektronisch berichtenverkeer met de overheid op de administratieve lasten van bedrijven
- 2008-21 Krisztian Balog (UVA)
People Search in the Enterprise
- 2008-22 Henk Koning (UU)
Communication of IT-Architecture

- 2008-23 Stefan Visscher (UU)
Bayesian network models for the management of ventilator-associated pneumonia
- 2008-24 Zharko Aleksovski (VU)
Using background knowledge in ontology matching
- 2008-25 Geert Jonker (UU)
Efficient and Equitable Exchange in Air Traffic Management Plan Repair using Spender-signed Currency
- 2008-26 Marijn Huijbregts (UT)
Segmentation, Diarization and Speech Transcription: Surprise Data Unraveled
- 2008-27 Hubert Vogten (OU)
Design and Implementation Strategies for IMS Learning Design
- 2008-28 Ildiko Flesch (RUN)
On the Use of Independence Relations in Bayesian Networks
- 2008-29 Dennis Reidsma (UT)
Annotations and Subjective Machines - Of Annotators, Embodied Agents, Users, and Other Humans
- 2008-30 Wouter van Atteveldt (VU)
Semantic Network Analysis: Techniques for Extracting, Representing and Querying Media Content
- 2008-31 Loes Braun (UM)
Pro-Active Medical Information Retrieval
- 2008-32 Trung H. Bui (UT)
Toward Affective Dialogue Management using Partially Observable Markov Decision Processes
- 2008-33 Frank Terpstra (UVA)
Scientific Workflow Design; theoretical and practical issues
- 2008-34 Jeroen de Knijf (UU)
Studies in Frequent Tree Mining
- 2008-35 Ben Torben Nielsen (UvT)
Dendritic morphologies: function shapes structure

====
2009
====

- 2009-01 Rasa Jurgelenaite (RUN)
Symmetric Causal Independence Models
- 2009-02 Willem Robert van Hage (VU)
Evaluating Ontology-Alignment Techniques
- 2009-03 Hans Stol (UvT)
A Framework for Evidence-based Policy Making Using IT
- 2009-04 Josephine Nabukenya (RUN)
Improving the Quality of Organisational Policy Making using Collaboration Engineering
- 2009-05 Sietse Overbeek (RUN)
Bridging Supply and Demand for Knowledge Intensive Tasks - Based on Knowledge, Cognition, and Quality
- 2009-06 Muhammad Subianto (UU)
Understanding Classification
- 2009-07 Ronald Poppe (UT)

Discriminative Vision-Based Recovery and Recognition of Human Motion

- 2009-08 Volker Nannen (VU)
Evolutionary Agent-Based Policy Analysis in Dynamic Environments
- 2009-09 Benjamin Kanagwa (RUN)
Design, Discovery and Construction of Service-oriented Systems
- 2009-10 Jan Wielemaker (UVA)
Logic programming for knowledge-intensive interactive applications
- 2009-11 Alexander Boer (UVA)
Legal Theory, Sources of Law and the Semantic Web
- 2009-12 Peter Massuthe (TUE, Humboldt-Universitaet zu Berlin)
perating Guidelines for Services
- 2009-13 Steven de Jong (UM)
Fairness in Multi-Agent Systems
- 2009-14 Maksym Korotkiy (VU)
From ontology-enabled services to service-enabled ontologies (making ontologies work in e-science with ONTO-SOA)
- 2009-15 Rinke Hoekstra (UVA)
Ontology Representation - Design Patterns and Ontologies that Make Sense
- 2009-16 Fritz Reul (UvT)
New Architectures in Computer Chess
- 2009-17 Laurens van der Maaten (UvT)
Feature Extraction from Visual Data
- 2009-18 Fabian Groffen (CWI)
Armada, An Evolving Database System
- 2009-19 Valentin Robu (CWI)
Modeling Preferences, Strategic Reasoning and Collaboration in Agent-Mediated Electronic Markets
- 2009-20 Bob van der Vecht (UU)
Armada, An Evolving Database System