UNIVERSITY OF SOUTHAMPTON
FACULTY OF PHYSICAL AND APPLIED SCIENCES
Electronics and Computer Science

Mechanism Design for Information Aggregation within the Smart Grid

by

Harry Thomas Rose

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

February 2013

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF PHYSICAL AND APPLIED SCIENCES
Electronics and Computer Science

Doctor of Philosophy

MECHANISM DESIGN FOR INFORMATION AGGREGATION WITHIN THE SMART GRID

by Harry Thomas Rose

The introduction of a smart electricity grid enables a greater amount of information exchange between consumers and their suppliers. This can be exploited by novel aggregation services to save money by more optimally purchasing electricity for those consumers. Now, if the aggregation service pays consumers for said information, then both parties could benefit. However, any such payment mechanism must be carefully designed to encourage the customers (say, home-owners) to exert effort in gathering information and then to truthfully report it to the aggregator. This work develops a model of the *information aggregation problem* where each home has an autonomous *home agent*, which acts on its behalf to gather information and report it to the *aggregation agent*. The aggregator has its own *historical consumption information* for each house under its service, so it can make an imprecise estimate of the future aggregate consumption of the houses for which it is responsible. However, it uses the information sent by the home agents in order to make a more precise estimate and, in return, gives each home agent a reward whose amount is determined by the payment mechanism in use by the aggregator. There are three desirable properties of a mechanism that this work considers: *budget balance* (the aggregator does not reward the agents more than it saves), *incentive compatibility* (agents are encouraged to report truthfully), and finally *individual rationality* (the payments to the home agents must outweigh their incurred costs). In this thesis, mechanism design is used to develop and analyse *two mechanisms*. The first, named the *uniform mechanism*, divides the savings made by the aggregator equally among the houses. This is both Nash incentive compatible, strongly budget balanced and individually rational. However, the agents' rewards are not fair insofar as each agent is rewarded equally irrespective of that agent's actual contribution to the savings. This results in a smaller incentive for agents to produce precise reports. Moreover, it encourages undesirable behaviour from agents who are able to make the loads placed upon the grid more volatile such that they are harder to predict. To resolve these issues, a novel scoring rule-based mechanism named *sum of others' plus max* is developed,

which uses the *spherical scoring rule* to more fairly distribute rewards to agents based on the accuracy and precision of their individual reports. This mechanism is weakly budget balanced, dominant strategy incentive compatible and individually rational. Moreover, it encourages agents to make their loads *less* volatile, such that they are *more* predictable. This has obvious advantages to the electricity grid. For example, the amount of spinning reserve generation can be reduced, reducing the carbon output of the grid and the cost per unit of electricity. This work makes use of both theoretical and empirical analysis in order to evaluate the aforementioned mechanisms. Theoretical analysis is used in order to prove budget balance, individual rationality and incentive compatibility. However, theoretical evaluation of the equilibrium strategies within each of the mechanisms quickly becomes intractable. Consequently, empirical evaluation is used to further analyse the properties of the mechanisms. This analysis is first performed in an environment in which agents are able to manipulate their reports. However, further analysis is provided which shows the behaviour of the agents when they are able to make themselves harder to predict. Such a scenario has thus far not been discussed within mechanism design literature. Empirical analysis shows the sum of others' plus max mechanism to provide greater incentives for agents to make precise predictions. Furthermore, as a result of this, the aggregator increases its utility through implementing the sum of others' plus max mechanism over the uniform mechanism and over implementing no mechanism. Moreover, in settings which allow agents to manipulate the volatility of their loads, it is shown that the uniform mechanism causes the aggregator to *lose* utility in comparison to using no mechanism, whereas in comparison to no mechanism, the sum of others' plus max mechanism causes an increase in utility to the aggregator.

Contents

List of Tables

List of Figures

Declaration of Authorship

I, Harry Thomas Rose, declare that the thesis entitled "Mechanism Design for Information Aggregation within the Smart Grid" and the work presented in the thesis are both my own, and have been generated by me as the result of my own original research. I confirm that:

- this work was done wholly or mainly while in candidature for a research degree at this University;
- where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
- where I have consulted the published work of others, this is always clearly attributed;
- where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
- I have acknowledged all main sources of help;
- where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
- parts of this work have been published as:

Rose, H., Rogers, A. and Gerding, E. H. (2011) Mechanism Design for Aggregated Demand Prediction in the Smart Grid. *In: AAAI Workshop on Artificial Intelligence and Smarter Living: The Conquest of Complexity*, 7-8 August, San Francisco.

Rose, H., Rogers, A. and Gerding, E. H. (2012) A Scoring Rule-Based Mechanism for Aggregate Demand Prediction in the Smart Grid. *In, The 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, Valencia, Spain

Signed: ______________________ Date: ______________________

Acknowledgements

First and foremost, I'd like to thank my supervisors, Dr. Alex Rogers and Dr. Enrico H. Gerding. Their guidance throughout my PhD and their support and encouragement was of immeasurable value to me and for that I am truly thankful. In addition, I'd like to thank the Engineering and Physical Sciences Research Council from whom I received the funding for my PhD via the ALADDIN DTA[1]. Furthermore, I'd like to thank all of the members of the Agents, Interaction and Complexity group at the University of Southampton who were of great support throughout my PhD.

I'd like to thank my parents, Amanda and Grenville Rose, who have always impressed upon me the importance of education and have always supported me through the highs and lows of academic life. In addition, I'd like to thank my birth father, Kevin Ormerod, who I'm sure is responsible for sparking my interest in computers, electronics and technology. Furthermore, I'd like to thank Neil Humphries, who provided me with valuable tuition in maths during high school. His extra tuition allowed me to gain the necessary grades for admission to college and ultimately, university. Mathematics is a tool I've used extensively throughout my undergraduate and postgraduate career and so I truly am thankful for his help.

Finally, I'd like to say thank you to all of the friends I've made at Southampton, both during my undergraduate degree and also during my PhD.

[1]EPSRC reference: EP/C548051/1

Nomenclature

N	The set of home agents $N = \{1, \cdots, n\}$.
x_i	Agent i's real distribution.
$\hat{x}_i$	Agent i's reported distribution.
$x_{a,i}$	The aggregator's prior belief of agent i's consumption.
μ_i	The mean of agent i's real distribution.
θ_i	The precision of agent i's real distribution.
$\hat{\mu}_i$	The mean of agent i's reported distribution.
$\hat{\theta}_i$	The precision of agent i's reported distribution.
$\mu_{a,i}$	The mean of the aggregator's prior information regarding agent i's event.
$\theta_{a,i}$	The precision of the aggregator's prior information regarding agent i's event.
$\ddot{\omega}_i$	The unmanipulated value of the event which agent i is trying to predict.
ω_i	The realised (possibly manipulated) value of the event which agent i is trying to predict.
$\boldsymbol{x}$	The vector of all of the agents' reported distributions.
$\boldsymbol{x}_a$	The vector of the aggregator's priors for each agent.
$\boldsymbol{\omega}$	The vector of the agents' realised events.
f	The forward market price.
δ^b	The system buy price delta.
δ^s	The system sell price delta.
f_r	The retail price per unit of electricity.
γ_i	The load variance manipulation factor for agent i.
$\alpha_{\gamma,i}$	The load variance manipulation cost coefficient for agent i.
$\alpha_{\theta,i}$	The precision cost coefficient for agent i.
$C_\gamma(\gamma_i, \alpha_{\gamma,i})$	The cost of performing load variance manipulation.
$C_\theta(\theta_i, \alpha_{\theta,i})$	The cost of generating a report of precision θ_i.
θ_{max}	The maximum precision accepted by the sum of others' plus max mechanism.
$\chi(\boldsymbol{x})$	The amount of electricity the aggregator must buy to minimise its exposure to the balancing markets given a vector of beliefs, $\boldsymbol{x}$.
χ, χ_a	Shorthand for $\chi(\boldsymbol{x})$ and $\chi(\boldsymbol{x}_a)$ respectively.

$\kappa(\chi, \omega)$	The cost of electricity given that χ units were initially purchased but ω units were consumed.
$\Delta(\boldsymbol{x}, \boldsymbol{x}_a, \boldsymbol{\omega})$	The savings made by using the agents' reports instead of the aggregator's priors.
λ	The fraction of the savings that should be distributed to the home agents as rewards.
$\mathbb{P}$	The set of probability distributions.
$\mathbb{R}$	The set of real numbers.
$P_i(\cdot)$	The reward function for agent i.
$P_\omega(\omega_i, \ddot{\omega}_i)$	The utility gained by an agent from consuming ω_i units of electricity when it actually requires $\ddot{\omega}_i$ units.
$U_i(\cdot)$	Agent i's utility function.
$U_a(\cdot)$	The aggregator's utility function.
$\bar{U}_i(\cdot)$	Agent i's expected utility function. The over-bar syntax is used throughout the thesis to denote the 'expected' variant of a function (e.g. $\bar{\Delta}(\cdot)$ refers to the *expected savings* function).

To my grandfather, Alan Norris.

Chapter 1

Introduction

Over the past decade, human-induced climate change has been thrust into the media spotlight and numerous agreements have been signed in an international effort to cut global greenhouse emissions. The current global targets, agreed to during the G8 summit in 2008, are to cut emissions by 50% compared to 1990 levels by the year 2050 (Department of Energy and Climate Change, 2009b), and as part of this effort, the UK has committed to cutting emissions by 20% by 2020, and by 80% by 2050. It is estimated that electricity generation contributed about 32% of the UK's total carbon emissions in 2007 (Department of Energy and Climate Change, 2009a, p. 125), and with the change towards using electrically powered vehicles and the electrification of heating through the use of heat pumps, the modernisation of the electricity grid is a key element of the UK Low Carbon Transition Plan (Department of Energy and Climate Change, 2009b).

In order to meet these carbon targets, the use of renewable sources of generation will become increasingly prevalent within the grid. However, these sources tend to be much more stochastic in nature compared to traditional sources of generation such as coal or gas (Varaiya et al., 2011). For example consider wind power, in the absence of storage, the electrical output of a wind farm is roughly proportional to the wind speed. Consequently, as wind-speed at any particular instant is difficult to predict, so is generation output from wind farms. Predictability within the grid is important because there are strict constraints on electricity networks stating that supply must *always* match demand. If these constraints are not met, the stability of the grid as a whole is threatened, with increased probability of transmission failure, brownouts (in which the voltage at the consumer end is considerably lower than required) or even blackouts. Thus, the unpredictability of renewable generators such as wind poses a considerable risk to the grid. To mitigate this risk, generation is planned ahead of time. Generators must declare the amount of electricity they are able to generate ahead of time, and these offers are matched with bids for electricity from consumers. However, any imbalance between what a generator commits to produce and the amount it actually produces must be rectified in real time by the grid. This is done by way of a *balancing mechanism* in which the generators who under- or over-supplied are charged costly fees for their imbalance. Therefore, there is clearly a need for generators to mitigate the risk they incur through their unpredictability.

Similarly, the demand placed on the grid is itself stochastic. Electricity companies must purchase sufficient supply for their customers in advance, and to do this they make use of statistical analysis plus some amount of 'tweaking'. If their customers do not consume as expected, thereby causing a surplus or deficit in demand, the aforementioned balancing mechanism is used to rectify the imbalance, and the electricity company is charged accordingly.

In both of these cases, customers might mitigate their risk by banding together under *aggregation services*. For example, in the case of wind farms, a number of geographically diverse wind farms might join an aggregation service, which then trades electricity on their behalf. Doing this reduces the risk of generators being unable to supply their agreed amount, as if one wind farm is unable to provide electricity, another in a remote location might be able to. Aggregation services will become particularly important due to the increasing use of distributed small-scale generators. Individually, it is not feasible for small generators to trade within the grid both due to the negligible amount of power these generators can provide in comparison to the amount being traded within the electricity markets, and due to computational constraints on the markets themselves (Pudjianto et al., 2007). Consequently, to be effective within the electricity markets, these small generators must be aggregated to form a *virtual power plant* in which an aggregation service trades the capacity of the individual generators on the grid on their behalf.

Further aggregation of suppliers is used to provide short term supplies of energy in order to supply peak or unexpected demand. Such services already exist in the UK thanks to the National Grid's (the system operator of the UK) short term operating reserve (STOR) program – a balancing mechanism that makes use of demand response in order to correct imbalance within the grid. Under demand response, loads placed by electricity consumers are curtailed either automatically or through the use of price cues. Two such examples are Flexitricity[1] and UK Power Reserve[2], which aggregate surplus generation capacity from small generators such that when a deficit of supply arises upon the grid, the aggregated generation capacity can be deployed.

On the side of the consumers, aggregators so far have concentrated on *demand response*. That is, they aggregate customers' ability to shed loads in the event of over-consumption on the grid. There are quite a number of these aggregation services currently registered

[1] `http://flexitricity.com`
[2] `http://ukpowerreserve.com`

with the National Grid – KiWiPower[3], Open Energi[4] and EnerNOC[5] to name just a few. However, these services tend to focus upon large consumers of electricity as it is them who are often able to shed the most load.

Home consumers can also benefit from aggregation by pooling their demand in much the same way as small generators pool their supplies. Aggregating consumers has the intrinsic benefit of reducing risk because it increases the likelihood of two or more prediction errors cancelling with one-another. However, further reduction of risk can be achieved through aggregating more specific information regarding each consumer – a problem on which this thesis focuses. This is made possible by the development of the 'traditional' electricity grid into a *smart grid*. The vision of the smart grid is a grid not only of electricity, but also one of information (U.S. Department of Energy, 2003; Electricity Networks Strategy Group, 2009). One key element in the vision of a smart grid is *smart meters*, which will be used to allow companies to send real time price signals to users to influence the demand that these consumers place on the grid. Such signals could be used, for example, to discourage users from placing a load on the grid when demand is already high (as in the case of real-time pricing), or when the carbon output is high (Ramchurn et al., 2011). However, crucially, this flow of information can be *bidirectional*, allowing consumers to send information towards the suppliers. The bidirectional flow of information offers opportunities for innovative, novel aggregation services to act as intermediaries between consumers and the electricity markets. This reduces the computational burden upon the markets as groups of consumers appear as a single entity.

It is this *demand prediction* problem on which this thesis focuses. However, the work within this thesis is equally applicable to supply prediction, or even the prediction of a customer's ability to shed loads. In fact, the model presented in Chapter 3 is an extension of the news vendor problem, in which a centre must predict the future demand of a perishable product such that it can purchase the correct amount. Importantly, in the newsvendor problem, the centre incurs cost penalties if it under or over-predicts. Consequently, the work in this thesis will in fact generalise to any such news vendor problem, or indeed, any problem in which a centre must perform some task whose expected cost is dependent upon the precision of the prediction the centre has about a future outcome. A more detailed discussion of the newsvendor problem is presented in Section 2.3.2.

The model presented in this thesis, which is more fully described and formally defined in Chapter 3, defines an aggregator and numerous homes represented by agents. Each home agent must gather information which can be used to precisely predict that home's future consumption. The aggregator then uses that information in order to predict the

[3] `http://kiwipowered.com`
[4] `http://openenergi.com`
[5] `http://enernoc.com`

cumulative demand of all homes under its service. Of course, the aggregator has some prior information regarding each home's expected consumption. This is the information it would have used had no additional information been available from the agents. As such, the aggregator is able to make some savings in using the agents' information over its own. However, the home agents incur a cost in gathering information for the aggregator, and therefore the aggregator must appropriately compensate them. Payments must be carefully designed as agents will strategise over their actions in order to perform actions that maximise their expected utility. Therefore, naive payments might cause agents to behave in a manner detrimental to the aggregator. It is with the design of these payments that this thesis is concerned.

1.1 Research Requirements

In line with the above discussion, the aim in this thesis is to develop an interaction mechanism between individual homes and the aggregator which includes a payment scheme that rewards individuals for sharing information with the aggregation service. In so doing, from the scenario in the previous section, the following requirements must be taken into account:

I **Automated:** The smart meter should use a piece of software to periodically collect information from the user. This should be easy and quick to use, requiring minimal effort from the user. Other than this interaction with the user, all other behaviour with regards to processing the information and communicating it with the supplier should be completely autonomous.

II **Encourage participation:** It must be profitable for each of the parties (the aggregator and the consumers) to purchase and sell electricity through this mechanism. In a branch of economics called *game theory*, individuals are assumed to be *rational*, meaning they will only perform actions that maximise their expected payoff (John von Neumann and Morgenstern, 1944). Consequently, if an individual is better off through not participating in the mechanism, it won't participate. Thus, in order to encourage individuals to use the mechanism, the savings made through the use of the system must outweigh the costs incurred by the parties from participating in it.

III **Budget balance:** The mechanism must not distribute more money to the home agents than is saved by the aggregator in using their information. That is, when savings are positive, the aggregator should never make a loss. Note that, assuming the agents truthfully report more precise information to the aggregator, the *expected* savings will be positive. However, there could be outlying occasions in which the information reported by the agents is less *accurate* than the information held by the aggregator. That is, on rare occasions, the aggregator might purchase an amount of electricity further from the actual amount consumed by the agents when using the agents information over its own. In such situations, weak budget balance will mean that the aggregator will make a loss (as the budget – the savings – is negative). Consequently, weak budget balance relates to incentive compatibility of the aggregator insofar as a weakly budget balanced mechanism is ex ante incentive compatible to the aggregator (i.e. in expectation the aggregator will make positive utility).

IV **Make effective use of information:** The aggregator will use the information reported to it by agents in order to optimally purchase electricity for those agents. Agents should therefore use all information that is available to them in order to make their predictions.

V **Increase the accuracy of predictions:** Consumers should be incentivised to produce accurate predictions of their future consumptions.

VI **Discourage the misreporting of consumers' predictions:** The aggregator must use the agents' information in order to optimally purchase electricity for them. Consequently, consumers should be encouraged to honestly report their beliefs over what their consumption will be, as if the aggregator tries to optimise over false information, the purchasing strategy it adopts is likely to be sub-optimal. Truth-telling should be incentivised both in the short and the long term, i.e. a consumer should not have a greater reward in the long term (say over a week) if they lie about their beliefs tomorrow, but report honestly for the next 6 days.

VII **Encourage consumers to report their confidence in their predictions:** In order for the aggregator to judge the accuracy of the reported predictions, consumers should also be encouraged to report their uncertainty in their beliefs. That is, when reporting their prediction, rather than just reporting their expectation, they should report their belief in the form of a probability distribution over their consumption. In a scenario whereby the cost of punishment for over- and underestimating future consumption is not symmetric, simply reporting a weighted mean of your estimated consumption is not sufficient for the aggregator to optimally purchase electricity. The aggregator must also be aware of the agent's confidence in its estimated mean – the variance of the agent's reported belief – in order to purchase the amount that, given the agents' predicted error, minimises the expected cost of the agents' electricity.

VIII **Fairness:** The rewards agents receive must reflect the effort they put into generating their information. We assume that an agent's effort is reflected in the precision of its estimate. That is, if an agent puts a large amount of effort into producing a precise estimate, it should be rewarded more greatly than an agent who put in less of an effort.

IX **Discourage inefficient use of electricity:** In order to reduce the carbon output of the grid, electricity must be used efficiently by consumers. The system should penalise homes that wilfully waste electricity.

X **Reduce the variance of demand on the grid:** The grid is rarely operating at maximum capacity, and variance in demand is often supplied for using standby generators that output more carbon than the generators used to supply the base load of the grid. By reducing the variance of demand on the grid, less standby generation capacity is required. This reduces development and maintenance costs, and allows cleaner sources to supply a greater percentage of the total electrical demand as the base line accounts for a greater amount of the total demand on the grid.

XI **Compatible with existing markets:** The model should not require a redesign of the current electricity markets. The redesign of said markets without thorough testing can lead to disastrous consequences (Borenstein, 2002). The model should make use of the markets that already exist to trade electricity, allowing for easier implementation and more immediate applicability.

The key requirement for the aggregator – who will ultimately implement this system – is that, it will reduce its expenditure through implementing this mechanism over simply using statistical methods over historical data.

1.2 Research Challenges

Given the requirements stated in Section 1.1, some key technologies and tools that will be used throughout this thesis are identified.

Requirement I states that autonomy is required by the system, which suggests the use of software agents. These are pieces of software which represent individual actors within a system and exhibit *autonomous, goal-driven behaviour* (Wooldridge and Jennings, 1995; Castelfranchi, 1995). That is, they perform actions that they believe will bring them closer to achieving some goal which they have been set. Agents are completely autonomous and are able to communicate with one another, allowing the exchange of information (for example, between homes and the aggregation service). Software agents are typically assumed to be *self-interested* and *rational* and, as a result they only perform tasks which they believe will benefit themselves. A branch of economics named *game theory* (Fudenberg, 1991) is used to understand the strategic behaviours each agent exhibits.

Furthermore, the information about each home's predicted consumption is private to that individual home's agent and it is this information that the aggregator would like to elicit from said agents. Private, in this sense, refers to the fact that only one agent knows the consumption for its home (as opposed to the information being confidential). This privacy occurs as a result of each agent gathering information independently to construct an estimate it will report to the aggregator. If an agent's information is its *type*, the problem can be thought of, in game theoretic terms, as a game of incomplete information. A branch of game theory called *mechanism design* (Mas-Colell et al., 1995, ch. 23) is particularly relevant to this problem, as it focusses on games in which agents' types (in the scenario from this thesis, their predictions and costs) are private. In more detail, mechanism design is a game theoretic tool with which a game can be designed such that specific behaviours are elicited from the participating agents. In particular, it introduces two key concepts – individual rationality and incentive compatibility. The concept of individual rationality states that an agent always seeks to maximise its utility and a mechanism that is said to be individually rational has the property that rational agents will participate in it. A mechanism is said to be incentive compatible when the reward received by an agent is maximised when that agent truthfully reports its beliefs. An example of an incentive compatible mechanism through which agents' private types are elicited is the Vickrey auction (Vickrey, 1961). This is a second price, sealed bid auction in which a good is being sold to the agent who values that item the most. However, the value an agent assigns to an item is private to that agent (i.e. it is not known by the auctioneer or the other bidders). Therefore, the auctioneer must devise a method of eliciting each agent's attributed value. The Vickrey auction does this by decoupling the agent's individual bid from what that agent pays. Specifically, an agent pays the next highest price bid (if they win). This decouples the amount the agent pays from the report it gives, thereby inducing incentive compatibility in the mechanism. For example, consider an agent bidding in a Vickrey auction. If it misreports its value (i.e. it bids a value other than its own true value), it will either decrease its probability of winning the auction in the case of reporting a lower value, or increase the probability of receiving a negative utility in the case of over reporting. Therefore, an agent's optimal strategy is to bid its true value.

The problem addressed by the Vickrey auction has interesting parallels to the scenario in which the aggregator tries to elicit the home agents' private information. However, in that scenario, there is no item with definite value being sold and agents are not bidding monetary values. Therefore, a typical auction mechanism cannot be applied to the problem. Instead, the setting of eliciting predictive information is considered and, in doing so, a closely related field of research, *scoring rules*, is discussed, which are often used in the evaluation of predictions to measure the quality of a prediction (Brier, 1950; Good, 1952; Savage, 1971; Gneiting and Raftery, 2007). One of the useful results of scoring rules is that of *propriety*, which is analogous to incentive compatibility in mechanism design. Specifically, if a scoring rule is strictly proper then an agent's utility is uniquely maximised when they truthfully reveal their types.

The majority of the work done so far on scoring rules has developed rules that take into account one distribution (the prediction). However, in order to address requirement IV, the aggregator must score the agent based on how useful the agent's report was to the aggregator. Scoring rules by Nau et al. (2007) take into account not only a distribution representing the prediction to be scored, but also a distribution that represents the aggregator's prior beliefs. This way, the score an agent achieves is based on the difference between its report and the aggregator's prior belief. However, this work cannot be directly applied to the problem discussed in this paper as it suffers from computational issues when the centre believes an event to be highly improbable.

An additional challenge to using scoring rules in the problem within this thesis is that mapping scores directly to payments in a manner is budget balanced (i.e. making the agents' payment proportional to the score they receive such that the aggregator doesn't make a loss) is non-trivial. This results from the fact that scoring rules are unbounded in continuous settings. Furthermore, while proper scoring rules and their affine transformations are incentive compatible, no guarantees are made with regards to other dimensions over which agents may strategise. For example, agents may be able to make themselves harder or easier to predict, as well as being able to manipulate their actual consumptions to match their earlier predictions. However, without additional analysis and design of the payment mechanisms, the behaviour of these agents is unknown. In fact, if an agent's payment is directly proportional to its score, it is likely that the agent will have incentive to waste energy in order to correctly meet its prediction, thereby violating Requirement IX.

Against this, the next section discusses the contributions of this thesis to the state of the art.

1.3 Research Contributions

The ultimate contribution of this work is the presentation and analysis of mechanisms for use within the smart grid, which allow an aggregation service to more optimally purchase electricity for a set of consumers. In doing this, mechanisms are presented which encourage the generation and truthful transmission of precise predictions of future consumptions by home agents. Moreover, this work presents mechanisms that encourage agents to reduce the variability of their loads and actively discourages agents from wasting energy in the event that they learn they will not consume as predicted. In doing this, the work contributes the *sum of others' plus max* mechanism to the literature, which makes direct use of scoring rules to assign payments to agents based upon a budget determined by their own reports in a budget balanced manner that is incentive compatible and encourages agents to make themselves more predictable by decreasing the volatility of their loads. In addition, this report discusses methods of removing the sum of others' plus max mechanism's incentive to waste.

As part of this discussion, this report discusses work presented in the following papers:

> Rose, H., Rogers, A. and Gerding, E. H. (2011) Mechanism Design for Aggregated Demand Prediction in the Smart Grid. *In: AAAI Workshop on Artificial Intelligence and Smarter Living: The Conquest of Complexity*, 7-8 August, San Francisco.
>
> Rose, H., Rogers, A. and Gerding, E. H. (2012) A Scoring Rule-Based Mechanism for Aggregate Demand Prediction in the Smart Grid. *In, The 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, Valencia, Spain

In more detail, the work detailed in this report contributes to the state of the art in the following ways:

- A new scoring rule-based mechanism named sum of others' plus max is presented, which can be applied to the information aggregation problem. The mechanism rewards agents using a budget determined by their own reports and takes into account both the agents' reports and the centre's prior information.
- The mechanism is the first to make direct use of agents' scores in continuous settings in order to make payments towards agents that are both incentive compatible and budget balanced.
- The mechanism is proven to be dominant strategy incentive compatible and ex ante weakly budget balanced.
- Using a computational approach to find equilibrium states, the mechanism is compared to a benchmark mechanism in which rewards are divided uniformly between agents, and another scenario in which agents' reports are not requested, and no reward mechanism is in place.
- Empirical analysis of the mechanisms under a smart-grid simulation is presented in order to show that the sum of others' plus max mechanism reduces the risk to the centre, and increases the centre's expected utility.
- It is shown that sum of others' plus max provides incentives for agents to *reduce the volatility* of their loads. Whereas the simpler, uniform mechanism provides the opposite incentives.
- It is shown that while under the sum of others' plus max mechanism, agents have incentive to waste electricity under certain conditions, the mechanism can be adjusted such that there is no such incentive.

With this in mind, the next section describes the structure of this thesis.

1.4 Thesis Structure

The remainder of this thesis is structured as follows:

Chapter 2 A review of the literature that is relevant to this thesis is presented. In more detail, this chapter presents a review of the theoretical papers behind proper scoring rules, as well as mechanism design papers relevant to the information elicitation problem. The chapter also discusses electricity markets and the simulation of consumers within the grid as well as the news vendor problem, and other models similar to the one presented here.

Chapter 3 A scenario which describes the problem discussed within this paper is presented. The scenario is that of a consumer within a smart grid setting purchasing electricity through an aggregation service, whose job is to interact with the electricity market on the consumer's behalf. The aggregator receives better rates when it precisely predicts the future consumption of the customer, and is able to make some imprecise prediction on its own. However, the customer is able to send more information regarding their future behaviours in order to allow the aggregator to more precisely predict their consumption. In return for this, the customer gets a rebate. The problem discussed within this thesis is *how this rebate is defined*. After the scenario is presented, the chapter goes on to formally present the problem and the model such that theoretical and empirical analysis may be performed on it.

Chapter 4 The information aggregation problem in which customers may misreport their information to the aggregator is discussed. This chapter discusses how payments can be made to the agents such that this behaviour is discouraged, and in so doing presents two such mechanisms. The payments discussed in this section are based upon the *savings* made by the aggregator in using the agents' information over its own. To this end, this chapter presents the *uniform* mechanism, which simply divides savings equally amongst agents and shows these payments to be Nash incentive compatible and budget balanced. Moreover, to obtain stricter incentive compatibility properties, the chapter presents the *sum of others' plus max* mechanism, which is the first mechanism to make direct use of scoring rules under a continuous setting in order to make payments to agents based upon the agents' reports in a way that is budget balanced and maintains incentive compatibility.

Chapter 5 The uniform and sum of others' plus max mechanisms are discussed under a setting that includes the ability for agents to *adjust the predictability of their events*. An agent adjusting its predictability directly affects the amount of information contained in the aggregator's information. When using budgets that are dependent on the savings made by the aggregator when using the agents' information over its own priors, it is possible for agents to increase these savings by making themselves harder to predict. This chapter performs empirical analysis of the uniform and sum of others' plus max mechanisms, discussing the equilibrium strategies of the agents and drawing comparisons to the results from Chapter 4. As well as manipulating the predictability of their consumptions, this chapter provides a theoretical analysis of agents' behaviours when it is assumed that agents are able to continuously gather information regarding their consumptions. Such an ability allows the agents to detect whether or not they will in fact consume as they had originally predicted. In the event of an agent not consuming as it had originally predicted, that agent is able to *fix* its consumption by 'wasting' electricity. This shows that while the uniform mechanism doesn't provide any incentive for agents to waste their electricity, the sum of others' plus max mechanism, on the other hand, does. A method of removing this incentive is discussed.

Chapter 6 Finally, this chapter provides a summary of this thesis as well as a discussion of the conclusions that can be drawn from the results in the preceding chapters. In addition, future extensions to this work are discussed, including further potential analysis to the mechanisms provided within this paper and potential extensions to the model presented within this thesis.

Chapter 2

Background

Given the introduction and research requirements discussed in the previous chapter, this chapter presents the background relevant to the work in this thesis. Additionally, it discusses the current state-of-the-art and the limitations that restrict its applicability to the problem described in Chapter 3. First, Section 2.1 presents a discussion of electricity markets and consumer behaviour. Next, Section 2.2 presents a discussion of smart grid technologies and how they enable the work within this thesis. Following that, Section 2.3 discusses a relevant problem from the field of inventory theory, namely that of the newsvendor problem. Section 2.4 provides a brief discussion of the aspects of game theory needed to understand this thesis. Then, Section 2.5 discusses relevant work in the fields of mechanism design and information elicitation, followed by a discussion of scoring rules in Section 2.6. Then, Section 2.7 discusses work surrounding coalition formation in which groups of consumers or producers form coalitions to buy or sell electricity as virtual power plants. Finally, Section 2.8 draws this chapter to a close by summarising the limitations of the current state of the art.

2.1 Electricity Markets

The goal of the British Electricity Trading Transmission Arrangements (BETTA), formerly the New Electricity Trading Arrangements (NETA), is to trade electricity as though it were any other resource, in as decentralised a manner as possible. This is done through bilateral trading between generators and suppliers in the forward markets, and then using a centralised balancing market to match supply with demand in real time.

The BETTA system splits the day into 48, half-hour trading slots in which electricity for that specific half-an-hour can be bought and sold. The participants in the BETTA system can be broadly split into two categories: generators and suppliers. In this case, generators are those who supply the grid with electricity, and suppliers those who supply end consumers with electricity. Electricity suppliers must buy enough electricity to meet their customers' demand at all times and, to do this, the electricity companies must estimate the amount of electricity their customers will require for each of these half hour slots. Ahead of time, electricity suppliers buy electricity in the form of long-term contracts and spot markets. However, an hour ahead of any actual time slot, gate closure occurs, meaning that electricity can no longer be purchased for this period from the forward markets. After gate closure, any excess or shortfall in the amount of electricity purchased for their customers, and the customers' demand must be balanced through buying or selling the difference in the *balancing market*.

During the balancing process, generators and electricity suppliers (from which the end consumers purchase their electricity) offer bids to increase or reduce the amount of electricity being generated or demanded in order to maintain the balance between supply and demand of electricity on the grid. In the balancing market, electricity is traded at two different prices: the system sell price (SSP), and the system buy price (SBP). These prices are calculated as weighted means of the bids to either reduce generation, or increase demand on the grid; or the bids to increase generation, or decrease the demand on the grid (Dettmer, 2002).

In order to reduce the costs incurred through balancing errors, consumers can be combined (aggregated) by novel aggregation services, which trade electricity for groups of houses. These services can act within the current electricity markets without need for any re-design of the existing trading mechanisms, thus they are readily applicable to modern-day electricity markets (Requirement XI). The reasoning behind this is that any redesign of markets must be thoroughly tested before their use. This is especially true for electricity markets, which are particularly sensitive as a result of supply always having to be matched with demand (Borenstein, 2002). Meanwhile, the BETTA markets are already well established and have proven to be fairly robust to global economic issues such as the 2001 Enron crisis (Hesmondhalgh, 2003). Therefore, this thesis develops a mechanism that *complements* existing markets by allowing communication between the aggregator and the homes of information that can be used by the aggregator in order to more precisely predict the homes' future consumptions of electricity. Insofar as the prediction of demand is *currently* concerned, thus far, electricity companies tend to use relatively simple statistical modelling of consumers and some amount of human input in order to predict future consumptions. There has been limited work on the modelling of loads of individuals and groups of houses (Herman and Kritzinger, 1993; Heunis and Herman, 2002; McQueen et al., 2004; Carpaneto and Chicco, 2008). Indeed, the distribution of loads seems to be highly dependent on the individual consumers themselves (Carpaneto and Chicco, 2008). However, gamma, beta, and log-normal distributions tend to be most often used to represent aggregate loads. Moreover, analysis by Carpaneto and Chicco (2008) show the gamma, log-normal and normal distributions to be the most appropriate for modelling residential loads. One key technology in enabling more precise predictions of future demand will be the intelligence built into the *smart grid*, as discussed in the next section.

2.2 The Smart Grid

Modernising the electrical grid will be a large area of work over the coming years as governments focus on attempting to increase the security and reliability of electricity supply, whilst also reducing the carbon emissions of electricity production and consumption (U.S. Department of Energy, 2003; Electricity Networks Strategy Group, 2009; Department of Energy and Climate Change, 2009b). One large area of research is based around implementing a *smart grid*, which comprises of nodes that are able to sense and act upon various variables, as well as communicate this information to one another. One such node is a smart meter, which is able to monitor and reflect the amount of electricity being consumed in real time. It then makes this information available both to the consumer and the electricity supplier, and can also receive pricing information in real time to allow real time pricing to be implemented, which is discussed next.

Various forms of feedback can be used to relay information to consumers. The most effective form, in terms of percentage savings made by customers, is direct feedback in which real-time price data is displayed clearly to a user (Darby, 2001, 2006). By clearly displaying that the current price of electricity is more expensive than average, consumers can decide whether or not they are able to ‘make do’ without placing a load on the grid. For example, a washing machine is not necessarily likely to have any imminent deadlines for finishing, so can be turned on later, when electricity is cheaper. This is called demand side management (as demand is shaped around the supply of electricity) and contrasts with the system commonly used today whereby *supply* is constantly varied to match demand in real time (Electricity Networks Strategy Group, 2009). These demand side management services must make a prediction regarding their customers’ ability to shed loads. Moreover, prediction is required even in scenarios without demand side management. For cleaner fuels, such as nuclear power, it is costly to vary output levels in real time. However, through the use of accurate predictions of future demand, preparations can be made to slowly ramp them to a minimum base-load such that a large percentage of electricity being generated is done so with minimum carbon-output. When combined with demand side management, the carbon output of the grid should be greatly reduced.

To further reduce the carbon and monetary cost of electricity consumption, aggregation can be performed within the home. In the HomeBots system, each electrical device is represented by an agent whose task it is to minimise the cost of running that device (Ygge and Akkermans, 1996; Ygge et al., 1999). This is an example of how demand-side

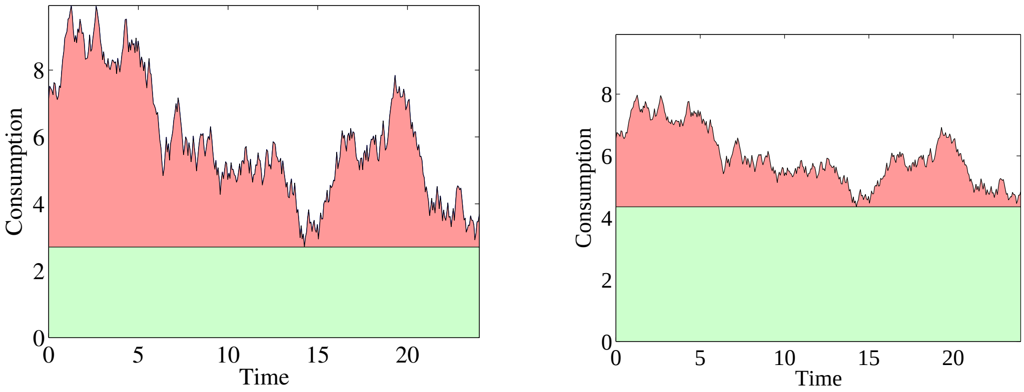

FIGURE 2.1: A pictorial example of how reducing variance in demand reduces the carbon output of the grid. The black line represents the demand being placed on the grid. The red area is electricity being supplied by more carbon-producing generators, whereas the green area is being supplied by more carbon-neutral generators.

management can be applied within a home in order to smooth the load placed on the grid by that home over time. This smoothing reduces carbon output. As a result, they tend only to be used to supply the guaranteed base load on the grid (as in Figure 2.1), whereas other, less carbon-neutral supplies are used to fill the gaps.

In the HomeBots system, load-smoothing is achieved by allowing device agents to bid for a time slot in which their device will consume power, and as part of this process, the system makes the important distinction between loads that are controllable and, those that are uncontrollable. Controllable loads are those which do not affect the resident of the property if they are interrupted or delayed, whereas uncontrollable loads are those that are provided on demand to the user, such as, a television or a desktop computer.

Although the work presented in this thesis does not look at demand management within houses (and indeed, the mechanisms designed in this work are aimed to be relevant to all aggregation services in the smart grid), the HomeBots system complements the work in this thesis, whereby tiers of aggregators are employed. For example, an aggregator could be placed at the home-level (the home agent), followed by an aggregator at the street or substation level. By distributing the aggregators in such a way, it is possible to achieve more precise predictions while reducing the computational load concentrated at any one point. In this system, individual appliances would know their load profiles and their function. The home agent could then use this information from the appliances, combined with other information that is available, in order to predict the load profile of the house.

The ability to accurately predict future consumptions and apply demand side management to the load on the grid helps to smooth the demand of electricity on the grid. However, it should be noted that some degree of flexibility is still required on the supply side of electricity generation. For example, simply relying on a large number of nuclear power stations to provide enough generation for the UK does not allow enough flexibility to cope with unexpected demand fluctuations. In the case that a sudden demand increase occurs, and there are insufficient 'flexible' power stations to supply it, brownouts[1] could occur, which can affect the stability of the grid as a whole. One of the contributing factors in the Californian energy markets in the early 2000s was the lack of flexibility in both the supply and demand side of the energy markets (Borenstein, 2002).

Furthermore, problems with inflexibility of supply are likely to be exacerbated in the future with the rise of *microgrids*. There are two main visions of the future electricity grid – a trans-European super-grid, and a grid of smaller (e.g. town-sized) smart grids each with their own means of storage and generation (European Commission and Directorate-General for Mobility and Transport, 2008; Galvin and Yeager, 2008). It is the latter, smaller grids, which are termed microgrids. With regards to these microgrids, the vision is of grids with many distributed sources of generation, mainly consisting of renewables. Renewables are inherently inflexible and therefore must be combined with electrical storage technologies. However, this inflexibility and lack of capacity gives rise to planning problems. For example, it must be ensured that at no point does consumption outweigh available generation capacity. In order to do this, accurate forecasts of consumptions must be used in order to source electricity from other connected microgrids for large demands that cannot be met solely by local generators. Therefore, demand-side management will be especially important in microgrids, as the supply side will be much less flexible in meeting fluctuating demand (Rahman et al., 2007). The work in this thesis is therefore particularly suited to micro-grids, with little flexibility in their generation capacity.

As far as electricity companies are concerned, deciding an optimal amount of electricity to purchase given a prediction of future demand is not a new problem. In fact, it is a similar problem to the problem faced by shops and commerce worldwide: the problem of inventory management, to which an entire field of theory is dedicated. The next section discusses aspects of inventory theory that are relevant to the smart grid model proposed in Chapter 3.

2.3 Inventory Theory

Inventory theory originates from the field of operations research. This theory deals with the problems faced by a centre who, given some demand for a product, must decide how much of that product to stock given certain costs (Hillier and Lieberman, 2005, ch. 18). The problem faced by the aggregator in the work presented in this thesis is an example of an inventory management problem, specifically, the newsvendor problem, discussed in Section 2.3.2. However, first, Section 2.3.1, discusses the relevant concepts and terminology from the field of inventory theory.

[1] A brownout is a drop in grid voltage that causes the grid to destabilise.

2.3.1 Terminology and Concepts of Inventory Theory

The field of inventory theory was developed to allow a centre to adequately cater for future demands. Broadly speaking, demand can be split into two classes. The first and simplest is *deterministic* or known demand. Under such a model, the centre knows the exact future demand that will occur for the product. The more interesting and, as far as this thesis is concerned, by far more relevant class of demand is that of *stochastic demand*. This model assumes that demand cannot be precisely known, but instead is randomly distributed. When modelling electricity markets, demand is assumed to be stochastic. Moreover, as discussed in Section 2.1, the demand placed on the grid tends to be best modelled by either log-normal distributions or gamma-based distributions.

When ordering an amount of product to supply future expected demand, the centre must take into account a number of costs. Those particularly relevant to the smart-grid model are: the *cost of ordering*, which corresponds to the cost per unit of electricity the centre incurs when purchasing electricity from the forward markets; the *shortage cost*, which is incurred when the demand is greater than the amount of product stocked (i.e. consumers have placed a load on the grid that is greater than the amount of electricity bought in the forward market); and the *salvage value*, which corresponds to the system buy price (i.e. the amount of money the centre can regain by selling excess electricity back to the grid).

Of course, dependent on the longevity of the product in question, the salvage value might never need to be taken into account by the centre. Products that are ever-lasting and for which there will forever be demand will never be sold at salvage value. These types of products are said to be *stable*. On the other hand, *perishable* products have a finite life-span, or will not be in demand after a certain period. Electricity is an example of a perishable product as any amount purchased must be either consumed by customers, or sold back to the grid (salvaged) in real time such that supply and demand are constantly equal. Models of such situations, in which all of a stocked product must be either sold or salvaged, are known as *single period* models. Consequently, the problem faced by the centre is one of a single-period model with stochastic demand, and perishable products. Such a problem has received a lot of attention among the inventory theory literature, and is commonly referred to as the *newsvendor problem*.

2.3.2 The Newsvendor Problem

The newsvendor problem is a wide-ranging problem in which a newsvendor must determine the number of newspapers to purchase given that he is unsure of what the demand for that paper will be.

The newsvendor incurs various costs dependent on how the quantity he orders matches the realised demand. Specifically, if the newsvendor orders too few newspapers, he incurs a *shortage cost*, which can represent either the loss in business as a result of the lack of newspapers, or conceivably the higher cost of procuring extra newspapers in short notice. Conversely, if he orders too many newspapers, he can scrap them in order to get some *salvage cost*, which may be zero (i.e. the papers have no value). In the newsvendor problem, a newsvendor is tasked with buying a quantity of newspapers for a stochastic future demand. Furthermore, the newspapers are said to be *perishable* as there is only demand for them on the day for which they were published. More formally, the newsvendor must choose a number of newspapers to purchase, S, given some probabilistic demand, D, whose probability density function is denoted by $f(x)$. Given this, the cost of purchase, c, the shortage cost, p, and the holding cost (the cost of storage minus the papers' scrap value), h, the newsvendor's expected cost is given by (Hillier and Lieberman, 2005):

$$C(S) = cS + \int_{S}^{\infty} p\,(x - S)\, f(x)\, \mathrm{d}x + \int_{0}^{S} h\,(S - x)\, f(x)\, \mathrm{d}x \tag{2.1}$$

That is, the newsvendor's expected cost is the cost of the initial purchase, plus the amount of money he expects to lose due to shortage, plus the cost he incurs for holding any unsold papers. This equation can be easily adapted to give the costs incurred by an aggregator within the smart grid, as is seen later in Chapter 3 (specifically, Equations 3.1 and 3.3). Furthermore, it lends credence to the argument about the wide applicability of the work within this thesis, as the model presented within Chapter 3 consists of the newsvendor problem imbued with the ability to elicit information from the newsvendor's customers; methods for this are discussed next.

2.4 Game Theory

Requirement I states that the system developed as part of this thesis must be autonomous. Given this, agent-based game theoretic modelling is used in order to simulate individual actors within the system. Moreover, an agent-based model is developed in which each house is represented by an autonomous, intelligent piece of software named an agent. Each agent is said to be self-interested, and chooses its actions based upon what it expects will maximise its reward for performing those actions (Wooldridge, 1999; Wooldridge and Jennings, 1995). Given this, the analysis within this thesis uses a branch of economics called game theory, which studies how such parties strategically interact with one another within a system (John von Neumann and Morgenstern, 1944). This section will introduce the core concepts of game theory, which are used throughout this thesis.

2.4.1 Concepts in Game Theory

Game theory can be divided into two broad sections – cooperative and non-cooperative game theory. The work in this thesis will look at non-cooperative game theory in order to model the interactions between the individual consumers and the aggregator. Therefore, discussion is restricted to non-cooperative game theory. First, games of complete information are discussed. This introduces key concepts of game theory including some concepts that are also used in studying games of incomplete information, which are discussed afterwards.

2.4.1.1 Games of Complete Information

In general, game theory models interactions between a set of agents (or players) $N = \{1, \cdots, n\}$ as a game, Γ. In this game, each agent, i, can perform some action $a_i \in A_i$, and has a set of preferences over all agents' actions, which are defined by that agent's utility function. Each agent chooses a strategy $\sigma_i \in \Sigma_i$, where Σ_i is the set of all probability distributions over all actions the agent can perform, $\mathbb{P}_{A_i}$, denoting the probability with which agent i will perform that action. The utility function is defined for each agent i as

$$U_i : \Sigma \to \mathbb{R},$$

where

$$\Sigma = \prod_{j \in N} \Sigma_j$$

Therefore, the set Σ holds all possible combinations of strategies for all agents where each combination is referred to as a *strategy profile*. An agent is said to prefer a strategy profile σ_i over σ_i' iff $U(\sigma_i) > U(\sigma_i')$, where, $\sigma_i, \sigma_i' \in \Sigma_i$, and a *self-interested rational* agent, i, is able to strategise over Σ_i in order to choose a strategy $\sigma_i \in \Sigma_i$ that maximises its utility. Formally the agent chooses σ_i given:

$$\sigma_i = \underset{\sigma_i' \in \Sigma_i}{\arg\max}\, U_i\left(\sigma_i' \times \Sigma_{-i}\right)$$

where

$$\Sigma_{-i} = \prod_{j \in N \setminus \{i\}} \Sigma_j$$

If a strategy, σ_i for agent i specifies that it will play a particular action with probability 1, that strategy is termed a *pure strategy*. Otherwise, that strategy is referred to as a *mixed strategy*. Given these elements, a strategic form game is defined as:

$$\Gamma = \langle N, \Sigma = \{\Sigma_i | \forall i \in N\}, U = \{U_i | \forall i \in N\} \rangle.$$

		P2	
		$\mathcal{D}$	$\mathcal{S}$
P1	$\mathcal{D}$	<u>-1</u>,<u>-1</u>	0,-2
	$\mathcal{S}$	-2,0	-0.5,-0.5

TABLE 2.1: Payoff matrix for the prisoners' dilemma where each agent can choose to either defect (i.e. implicate the other prisoner), $\mathcal{D}$, or to remain silent, $\mathcal{S}$.

Attention is now turned to a game named the prisoners' dilemma. This game is described by the *normal form* representation shown in Table 2.1, and consists of two agents $N = \{1, 2\}$, both arrested by the police. They are taken to separate rooms such that they are unable to collaborate with one-another and each prisoner is offered a choice – to remain silent, $\mathcal{S}$; or to defect and implicate the other prisoner, $\mathcal{D}$.[2] If one prisoner defects while the other remains silent, the silent prisoner will receive a two year sentence, while the defecting prisoner will go free. Alternatively, if both prisoners remain silent, the police will only be able to prosecute them on a lesser offence, resulting in each prisoner receiving a half-year sentence. Finally, if both agents defect, the police will prosecute both prisoners, and they will share the two year sentence, receiving one year each. This is described in Table 2.1.

This is a game of *complete information* and as such, each prisoner knows the possible strategies and preferences of the other prisoner. Given this, each prisoner can analyse the game to decide on the strategy he will play. From prisoner one's perspective, it can be seen that for both of prisoner two's strategies, agent one will maximise his utility by choosing to defect. Therefore, it is said that for prisoner one, $\mathcal{D}$ is a *dominant strategy*, as it maximises his utility regardless of his opponent's strategy. Formally, for an agent, i, with a strategy space of Σ_i, a strategy $\sigma_i \in \Sigma_i$ is dominant if

$$U_i\left(\sigma_i\right) \geq U_i\left(\sigma_i'\right), \forall\, \sigma_i' \in \Sigma_i \setminus \{\sigma_i\}$$

and that the strategy is strictly dominant if the inequality is strict. Looking from prisoner two's perspective, it can be seen that the same holds – defecting is a dominant strategy. Thus, assuming both prisoners are rational, they will therefore always choose to implicate the other prisoner (defect). It is said that the strategy $(\mathcal{D}, \mathcal{D})$ is a *dominant strategy equilibrium* as no agent has an incentive to deviate (to choose to play a different strategy) *regardless* of the other agents' adopted strategies. Interestingly, both agents can gain a better utility by both remaining silent (receiving a utility of -0.5 each rather than -1). However, this is *not* a Nash equilibrium as from the $(\mathcal{S}, \mathcal{S})$ strategy profile each agent is able to unilaterally deviate to $\mathcal{D}$ in order to gain a higher reward (i.e. a utility of 0 rather than -0.5).

Consider now a different game named the battle of the sexes. This game consists of two agents, $N = \{1, 2\}$, who are unable to communicate with one-another and must meet

[2]Therefore, in this case $A = \{(\mathcal{D}, \mathcal{D}), (\mathcal{D}, \mathcal{S}), (\mathcal{S}, \mathcal{D}), (\mathcal{S}, \mathcal{S})\}$

		P2	
		$\mathcal{O}$	$\mathcal{R}$
P1	$\mathcal{O}$	<u>3</u>,<u>1</u>	0,0
	$\mathcal{R}$	0,0	<u>1</u>,<u>3</u>

Table 2.2: Payoff matrix for the battle of the sexes game where each agent can choose to attend either the opera, $\mathcal{O}$ or a rock concert, $\mathcal{R}$.

each other at one of two events – the opera, $\mathcal{O}$; or a rock concert, $\mathcal{R}$. Agent 1 prefers $\mathcal{O}$ whereas Agent 2 prefers $\mathcal{R}$. However, no matter which event they attend, they prefer to be together than to be apart. This game is described in Table 2.2.

Nash (1950, 1951) states that any finite game is guaranteed to have *at least* one mixed strategy *Nash equilibrium*. A Nash equilibrium is a strategy profile from which no one agent is incentivised to deviate. The game in Table 2.2 has *three* Nash equilibria – two *pure strategy* Nash equilibria and a single *mixed strategy* Nash equilibrium. The two pure strategy equilibria, which are underlined in Table 2.2, are $(\mathcal{O}, \mathcal{O})$ and $(\mathcal{R}, \mathcal{R})$.

To find the mixed strategy Nash equilibrium in this game, the agents' expected utilities must be analysed. Let $\sigma_1 = (\pi_1, 1 - \pi_1)$. That is, agent 1 will play $\mathcal{O}$ with probability π_1 and $\mathcal{R}$ with probability $1 - \pi_1$. In the same way, let agent 2's strategy be $\sigma_2 = (\pi_2, 1 - \pi_2)$. Agent 1 therefore gains the utility:

$$U_1(\mathcal{O}, \sigma_2) = U_1(\mathcal{O}, \mathcal{O}) \cdot \pi_2$$
$$U_1(\mathcal{R}, \sigma_2) = U_1(\mathcal{R}, \mathcal{R}) \cdot (1 - \pi_2)$$

In a mixed strategy equilibrium, each agent must be indifferent between the utilities of the two actions given the other agents' strategy, so $U_1(\mathcal{O}, \sigma_2) = U_1(\mathcal{R}, \sigma_2)$. Now, with this in mind and given the utilities in Table 2.2, agent 2 is able to compute the mixed strategy it should adopt:

$$3\pi_2 = 1 - \pi_2$$
$$\therefore \pi_2 = \frac{1}{4}$$

Therefore, agent 2 must play the strategy $\sigma_2 = (0.25, 0.75)$, and it can be similarly determined that agent 1 plays $\sigma_1 = (0.75, 0.25)$.

These games are termed games of *complete information*. That is, at the first point the agents calculate the strategies they will adopt, all agents know the strategies open to all other agents and the utilities obtained by all agents for every outcome. If this assumption is relaxed and it is said that the agents do not know their opponents' preferences at the first point at which they begin to strategise against one-another, the game becomes one of *incomplete information*, which is discussed next.

2.4.1.2 Games of Incomplete Information

The games that have been described so far are examples of games of complete information. In these games, agents know all information there is to know about the game before they strategise against each other. This section describes games in which agents hold some *private* information regarding the game – games of *incomplete information.* In such games, any private information an agent has at the beginning of the game is referred to as its *type* (Myerson, 1991).

In discussing games of incomplete information, Harsanyi (1967) describes three general cases in which agents can be ignorant of information given in the normal form representation of games given in the previous section:

1 Agents do not know the physical outcomes of the game.

2 Agents do not know either their own utility function or the utility functions of the other agents.

3 Agents do not know the strategy spaces of either themselves or of the other agents.

Furthermore, it is argued that the cases 1 and 3 can both be generalised to case 2 whereby agents are unaware of the utility functions of the other agents in the game. Therefore, games of incomplete information can be generally defined as games in which agents do not know the other agents' utility functions (Harsanyi, 1967).

In games of incomplete information (also referred to as Bayesian games), agents are unable to play against their opponents' types as they did in the games discussed earlier. Instead, they must play against probabilistic beliefs of what their opponents' types are. Bayesian games are therefore defined as follows (Myerson, 1991):

$$\Gamma^b = \langle N, \Sigma, T, \mathbb{P}, U \rangle$$

where T_i is the set of possible types of agent i and,

$$T = \prod_{i \in N} T_i$$

and where, $\mathbb{P}_i \in \mathbb{P}$, and $\mathbb{P}_i : T_{-i} \to \mathbb{R}$ is the probability agent i has assigned to agent $-i$ (i.e. agent i's opponent) being of a given type, $t_{-i} \in T_{-i}$, given its own type, t_i, such that:

$$0 \leq \mathbb{P}_i\left(\cdot\right) \leq 1$$

and

$$\sum_{t_j \in T_{-i}} \mathbb{P}_i\left(t_j | t_i\right) = 1$$

Given this information, agents strategise using their *expected utility*, which is given by the function:

$$\bar{U}_i\left(\sigma_i|t_i,T_{-i},\mathbb{P}_i\right)=\int_{t_j\in T_{-i}}\mathbb{P}_i\left(t_j|t_i\right)\cdot U_i\left(\sigma_i|t_i,t_j\right)$$

and agent i will choose a strategy σ_i such that:

$$\sigma_i=\arg\max_{\sigma_i'\in\Sigma_i}\bar{U}_i\left(\sigma_i'|t_i,T_{-i},\mathbb{P}_i\right)$$

where the overbar syntax of $\bar{U}(\cdot)$ is used throughout this thesis to denote the *expected* value of a function.

As in games with complete information, Bayesian games also have equilibria termed Bayes or Bayes-Nash equilibria (Harsanyi, 1968). These equilibrium positions occur when each agent, i, is playing its best response to every other agents' strategy given its own type, t_i, and its own beliefs $\mathbb{P}_i$ about the other agents types, T_{-i}.

A classic example of a Bayesian game is an auction in which an object is being sold to which the buyers attribute a certain private value determined by that agent's type. Consider a first price sealed bid auction between two agents $N=\{1,2\}$, each agent $i\in N$ attributes a value $t_i\in\mathbb{R}$ to the object, which that agent keeps private to itself. Each agent must write a given bid amount $a_i\in\mathbb{R}$ on a piece of paper, which he will pay to the auctioneer if, after both agents have submitted their bids, his bid is the highest. The utility function for each agent i is given by:

$$U_i\left(a_i|t_i\right)=\begin{cases}t_i-a_i, & \text{if } a_i>a_j,\ \forall j\in N\setminus\{i\}\\ 0, & \text{otherwise}\end{cases}$$

Let the item being sold be a rare, antique vase. The bidding agents can be of two types: a collector, $\mathcal{C}$, and a shopkeeper, $\mathcal{S}$, who is interested in reselling the vase at a profit. That is, $T_i=\{\mathcal{C},\mathcal{S}\}$. The two bidders never meet one-another and so do not know each other's types when they make their bids. However, they each know their own type t_i, and therefore agent 1 knows that $t_1=\mathcal{C}$ and agent 2 knows $t_2=\mathcal{S}$. Collectors do not wish to make profit on their purchases (as they do not plan to resell them), therefore the value that agent 1 has attributed to the vase is represented by the normal distribution $\mathcal{C}=\mathcal{N}(10,5)$. Shop keepers, on the other hand, must make a profit and therefore agent 2's value for the vase is $\mathcal{S}=\mathcal{N}(5,3)$. The strategy space each agent, i, is able to strategise over is $\Sigma_i=\mathbb{P}_{\mathbb{R}+}$, where $\sigma_i\in\Sigma_i$ is the amount that agent i will bid for the vase.

The agents know that $P_i\left(t_j=\mathcal{C}|t_i=\mathcal{C}\right)=0.2$ and $P_i\left(t_j=\mathcal{S}|t_i=\mathcal{S}\right)=0.6$, and therefore are able to strategise over the amount they bid for the item. If the agents were to assume that the other agents bid truthfully (which will soon be seen to be a false assumption),

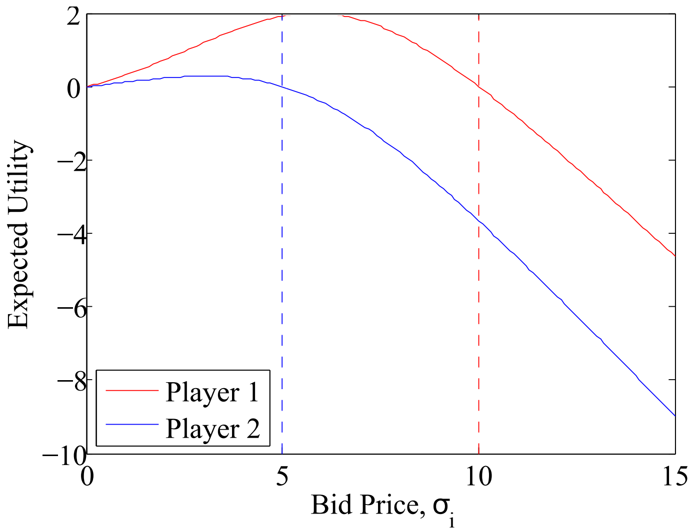

FIGURE 2.2: The expected utility for each agent assuming the other agent bids truthfully but when each agent is unaware of the other agents' realised type. The dashed line represents the agents' true mean value for the item.

the expected utilities for each agent i for each bid amount a_i can be seen in Figure 2.2. It can be seen that in this case, agent 1, whose type was $\mathcal{C}$ would in fact bid 5.8 as this maximises his expected utility, whereas agent 2, whose type is $\mathcal{S}$ would bid 3.1. The fact that both agents can maximise their expected utility by bidding a value other than their true private value shows this auction does not encourage truthful behaviour from agents, which is a concept that is discussed in more detail in Section 2.5.

2.5 Information Elicitation

Often, when designing markets and systems through which agents interact with one-another, the designer would like to guarantee certain properties of the system. For example, in order to make best use of the information available to the home agents, the aggregator must encourage the home agents to report that information truthfully (Requirement VI). Alternatively, a commonly used example is that of an auctioneer who would like to sell an object of unknown value. The auctioneer's goal is to efficiently allocate some goods to a set of bidders, which, in the case of a single item, means giving the good to the bidder who values the item the most. Therefore, the auctioneer must ascertain the value each bidder attributes to the item. In order to maximise the efficiency of the allocation he chooses, the auctioneer requires each bidder to report their *true* value (that way he can allocate the vase to the bidder who values it the most). The field of mechanism design provides the tools to construct Bayesian games which exhibit given properties, and as such is used extensively throughout this thesis. Therefore, this section introduces key concepts from mechanism design, and discusses relevant literature from the field.

In mechanism design, it is said that there is a Bayesian game containing a set of n agents, N, and for each agent, i, a utility function U_i, and a set of possible types T_i. Each agent, i, also has a defined set of actions, A_i, from which, when asked by the centre, he can choose to perform some action $a_i \in A_i$. Additionally, a mechanism defines an outcome space, Ω, which represents the set of actions the mechanism can perform. The mechanism chooses an outcome $\omega \in \Omega$ using a social choice function $F : A_1 \times \cdots \times A_n \rightarrow \Omega$. From this outcome, each agent obtains a certain value defined by the function $v_i : \Omega \rightarrow \mathbb{R}$. The mechanism designer must design a set of payments $P_i : A_1 \times \cdots \times A_n \rightarrow \mathbb{R}$, such that the desired set of actions are performed by the agents. Commonly, for example, a mechanism designer may want to design payments such that the agents truthfully reveal a piece of information to the centre.

Mechanism design can be thought of in three steps (Fudenberg, 1991):

1 The *principal* designs a mechanism that it expects will elicit specific behaviours from the agents.

2 The agents accept or reject the proposed mechanism. This occurs in one simultaneous move – i.e. agents have no information with regards to whether their peers have accepted/rejected – although they can speculate. At this point, agents who reject obtain a base utility.

3 The agents who accepted play the game specified by the mechanism and are rewarded with the base utility plus whatever utility they obtain from the result of the game.

Mechanism design introduces some key concepts such as *individual rationality*, being *budget balanced* and *incentive compatibility*, which are defined as follows:

Individual Rationality A mechanism exhibits individual rationality if it encourages *rational* agents to participate. That is, if an agent can choose between the strategies $\Sigma = \{a, b\}$, where a is to participate, and b is not to participate in the mechanism, a mechanism is individually rational if $\bar{U}(a) > \bar{U}(b)$.

Budget Balanced A mechanism that is budget balanced will never make a deficit. For example, if a mechanism must distribute a certain amount of money Ψ to n agents such that each agent, i, receives ψ_i, a mechanism is said to be *weakly* budget balanced if $\Psi \geq \sum_{\forall i} \psi_i$, and strongly budget balanced if the inequality is strict. Similarly, a mechanism can be *ex ante* or budget balanced in *expectation* if the above holds for expected payments – i.e. on average the above inequality holds.

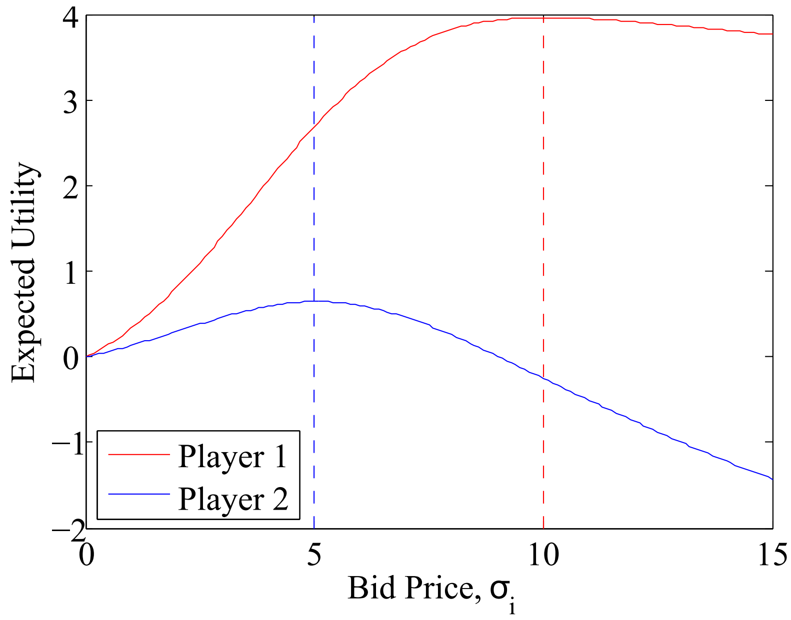

FIGURE 2.3: Agents' expected utilities versus their bid price under the Vickrey auction. It can be seen that the agents' utilities are maximised when they truthfully report their values to the auctioneer.

Incentive Compatibility A mechanism is said to be *incentive compatible* if there exists a Nash equilibrium for the strategy set whereby all agents report truthfully. Taking the auction example from the introduction to this section, the auctioneer would like all agents to truthfully report their value of the item being sold such that the item will be sold to the agent who values it the most. In this case, each agent i must report its true type t_i, and the mechanism is said to be incentive compatible if $\bar{U}_i\left(t_i\right) > \bar{U}_i\left(\hat{t}_i\right),\ \forall \hat{t}_i \in T_i \setminus \{t_i\}$.

There are two broad categories of mechanisms – *indirect* mechanisms whereby agents' messages are unrestricted; and *direct* mechanisms in which each agent, i's messages are restricted to reporting his set of types, T_i, to the centre. In mechanism design, often information elicitation is discussed in terms of agents revealing their *types*, rather than their private values, or their preferences. Moreover, an important result from the mechanism design literature is that of the *revelation principle* (Dasgupta et al., 1979; Myerson, 1979). This states that if an indirect mechanism can achieve an outcome in Bayes-Nash equilibria, there exists an equivalent, incentive compatible, *direct* mechanism in which that outcome is a Bayesian equilibrium. This way, attention can be restricted to only direct mechanisms (provided the possible types of each agent is common knowledge).

A solution to the auctioneer's problem from the beginning of this chapter was presented by Vickrey (1961). It was shown that by using a *second price*, *sealed bid* auction, agents would maximise their expected utility by truthfully reporting their types. This works by disassociating the agents' report from the amount it pays. It can be seen that each agent's expected utility is maximised when they report their true private value. At any bid price over their opponents, they still pay the same amount (the value of the next highest bidder's bid), but, if they bid lower than their true value, they reduce the probability that they are the highest bidder and therefore the winner. Of course, if they bid higher than their true private value, they improve their chances of winning the item, but also increase the probability of having to pay an amount higher than their true valuation of the object.

A similar application of mechanism design was used in Papakonstantinou et al. (2008, 2009), whereby an agent (whom is termed the centre) would like to obtain information from a set of sensory agents but does not know the cost of obtaining that information from those agents. Furthermore, the centre does not know the capabilities of the sensory agents and thus they might not be able to individually supply information of a high enough precision for the centre. In this case, in order to get an overall estimate with the required precision, θ_0, the centre will have to purchase information from multiple agents and then fuse those pieces of information. These works develop *two-stage* mechanisms, of which the first stage is discussed in this section and the second stage in Section 2.6.

In the case that the centre must only obtain information from one user (Papakonstantinou et al., 2008), it was demonstrated that a *reverse* second-price auction could be used in order to get the sensory agents to truthfully reveal their costs to the centre such that it can select the agent who will provide the information at the lowest cost. In this mechanism, the winning agent (the lowest bidder) would be paid the next highest amount over its own bid price, thus decoupling the agents' bid from its utility in the same way as the Vickrey auction. Papakonstantinou et al. (2009) then extends that work to address the issue whereby information must be sought from multiple agents. In this work, the centre asks n of the agents in N, where $n \geq 2$, to report their cost information so that it can select the m cheapest agents of those and discard the rest. The amount each of the m agents will be paid by the centre is the $m+1^{\text{th}}$ cost, i.e. the next highest price in the set of n agents after the highest price quoted in that group of m. That is, if there are $n = 5$ agents, and each agent $1, \cdots, 5$ quotes the prices $1, 2, 3, 4, 5$ respectively, if the centre chooses the $m = 3$ cheapest agents (those with costs $1, 2, 3$), each of the m agents will be paid 4 – the next lowest bid price outside of the m selected agents.

Note that although these auction-based mechanisms are commonly used in type-elicitation, they are unable to be applied to the smart grid scenario presented in Chapter 3 for numerous reasons. Foremost, the model here has only one agent who collects information regarding each home, whereas auction mechanisms rely on there being at least two. The mechanism in Papakonstantinou et al. (2009) also relies on fusing the information reported by the agents. However, this is not possible in the problem here as the information reported by the agents is *additive* – the aggregator would like to know the *total* consumption of all agents in the system.

Mechanism design has also been applied to situations whereby the information being elicited from the agents is not their costs, but rather the value of some variable to which they have access (for example, in the scenario presented in Section 3.1, this could be the consumption of the houses). Work has ranged from rating service providers, such as sellers on online auction sites (Miller et al., 2005; Gerding et al., 2009) to getting specific information regarding measured values (Chalkiadakis et al., 2011; Papakonstantinou, 2010; Zohar and Rosenschein, 2008, to name a few). Many of these solutions employ functions named scoring rules in order to achieve incentive compatibility, which are discussed in the next section.

2.6 Scoring Rules

In Section 1.2, it was stated that the mechanisms designed in this work must encourage agents to truthfully report their information about future energy consumptions (Requirement VI). Furthermore it was stated that the agents should report their confidence in their estimates (Requirement VII). Therefore, agents must report a *distribution*. Proper scoring rules take agents' probabilistic reports and a realised event and award them a score based on how closely their report predicted the realised event. These scores are maximised in expectation when agents report their belief truthfully and therefore, are often used in achieving incentive compatibility. Furthermore, they can be used to increase the fairness of rewards (Requirement VIII), as they can be used to individually reward agents based upon the precision and accuracy of their predictions (thereby also satisfying Requirement V). Scoring rules are functions that map an event and distribution to a real number (the score). More formally, a scoring rule is defined as $S : \mathbb{P} \times \Omega \to \mathbb{R}$ where Ω is the set of possible events and $\mathbb{P}$ is the set of all probability distributions over Ω. A scoring rule is (strictly) proper if it is maximised in expectation (only) when the agent truthfully reports its beliefs. That is, if an agent holds a belief $\boldsymbol{r}$, a scoring rule, S, is *proper* if:

$$\bar{S}\left(\boldsymbol{r}\right) \geq \bar{S}\left(\hat{\boldsymbol{r}}\right), \text{ where } \boldsymbol{r} \neq \hat{\boldsymbol{r}}$$

where $\bar{S}(\boldsymbol{r})$ is the expected score an agent will receive for reporting $\boldsymbol{r}$. Furthermore, S is termed *strictly proper* if the above inequality is strict.

2.6.1 Scoring Rules for Discrete Distributions

Scoring rules have been applied to information elicitation scenarios for decades. They were first used in meteorology, where it was noted that if forecasters were to report their predictions as probabilistic distributions and then to be rewarded through the application of a scoring rule to their report, they would maximise their expected reward only if they truthfully reported their forecast (Brier, 1950). The scoring rule used by

$\boldsymbol{r}$		Expected Score		
rain	$\neg$*rain*	Brier Score	Log Score	Spherical Score
0.1	0.9	-0.14	-1.64342	0.375467
0.2	0.8	0.08	-1.19355	0.460818
0.3	0.7	0.26	-0.949783	0.551487
0.4	0.6	0.4	-0.794651	0.637905
0.5	0.5	0.5	-0.693147	0.707107
0.6	0.4	0.56	-0.632465	0.748845
0.7	**0.3**	**0.58**	**-0.610864**	**0.761577**
0.8	0.2	0.56	-0.639032	0.75186
0.9	0.1	0.5	-0.764528	0.728848
1	0	0.4	$-\infty$	0.7

TABLE 2.3: The expected scores achieved by a forecaster when he believes $\boldsymbol{p}_{rain} = 0.7$ but reports $\boldsymbol{r}$. Note the peak score is at $\boldsymbol{r}_{rain} = 0.7$.

Brier has come to be known as the *Brier score*, and is defined as follows:

$$S(\boldsymbol{r}, \omega) = 2\boldsymbol{r}(\omega) - \sum_{\omega' \in \Omega} \boldsymbol{r}(\omega')^2 \tag{2.2}$$

where Ω refers to the set of events that can occur, $\boldsymbol{r}$ is the discrete probability distribution reported by the forecaster, and ω is the realised event that occurred. As an example of its use, imagine the scenario in which a town would like to know whether or not it will rain tomorrow. There are two possible events that can occur, $\Omega = \{rain, \neg rain\}$. The town enlists the help of their local weather forecaster, who they will pay an amount proportional to his Brier score for his report. The local weather forecaster predicts that tomorrow there is a 70% chance of *rain*. That is, $\boldsymbol{p} = \{0.7, 0.3\}$. However, he strategises over what he actually reports to the town, $\boldsymbol{r}$, in order to maximise his expected earnings, which he calculates using:

$$\bar{S}(\boldsymbol{r}|\boldsymbol{p}) = \sum_{\omega \in \Omega} \boldsymbol{p}(\omega) \cdot S(\boldsymbol{r}, \omega) \tag{2.3}$$

The scores the forecaster expects to receive for various values of his report, $\boldsymbol{r}$, are given in Table 2.3. It can be seen that the forecaster maximises his score when he reports $\boldsymbol{r} = \boldsymbol{p}$, and thus to maximise his reward, he should report his beliefs truthfully. The Brier score is an example of an *strictly proper* scoring rule. That is, in expectation an agent will maximise his score only when it reports its beliefs truthfully. A weaker notion is that of *proper* scoring rules, which are maximised when reporting truthfully, but can also be maximised by misreporting, making agents indifferent between reporting truthfully or misreporting.

The Brier score is synonymous with the quadratic scoring rule (which it will henceforth be called), and is one of three main scoring rules extensively studied throughout the literature. The spherical (Roby, 1965), and the logarithmic (Savage, 1971) scoring rule are the other two that have received much focus in the literature and are described by Equations 2.4 and 2.5 below.

$$\text{Log} : S(\boldsymbol{r}|\omega) = \log\left(\boldsymbol{r}\left(\omega\right)\right) \tag{2.4}$$

$$\text{Spherical} : S(\boldsymbol{r}|\omega) = \frac{\boldsymbol{r}\left(\omega\right)}{\sqrt{\boldsymbol{r} \cdot \boldsymbol{r}}} \tag{2.5}$$

2.6.1.1 The Characterisation of Discrete Scoring Rules

Discrete proper scoring rules were first formally characterised by Savage (1971), McCarthy (1956) and Schervish (1989). They state that given a convex cost function $G : \mathbb{P} \to \mathbb{R}$, where $\mathbb{P} \subset [0,1]^n$ is the set of all possible probability distributions over a set of n possible outcomes, Ω, a scoring rule, $S : \mathbb{P} \times \Omega \to \mathbb{R}$ is proper iff:

$$S(\boldsymbol{p}, \omega) = G(\boldsymbol{p}) - \left\langle G'(\boldsymbol{p}), \boldsymbol{p} \right\rangle + G'_i(\boldsymbol{p}) \; \forall i \in \{1, \cdots, n\} \tag{2.6}$$

where the term $\langle G'(\boldsymbol{p}), \boldsymbol{p}\rangle$ represents the dot product of the sub-derivative of G with respect to $\boldsymbol{p}$, $G'(\boldsymbol{p})$, and $\boldsymbol{p}$, and $G'_i(\boldsymbol{p})$ is the value of $G'(\boldsymbol{p})$ at event i.

It is shown in Theorem 2.1 that $\mathbb{P}$ is convex, and therefore, for G to be convex, it is sufficient to show that (R. Tyrrell Rockafellar, 1970):

$$G(\theta x + (1-\theta)y) \leq \theta G(x) + (1-\theta)G(y), \; \forall\, x, y \in \mathbb{P}, \; \forall\, \theta \in [0,1]$$

Theorem 2.1. *The set of discrete probability distributions over n mutually exclusive outcomes, $\mathbb{P} \subset [0,1]^n$ is convex.*

Proof. As stated in R. Tyrrell Rockafellar (1970), a set, G, is convex if

$$\begin{gathered} \forall\, \theta \in [0,1] \\ x \in G \wedge y \in G \Rightarrow (\theta x + (1-\theta)\, y) \in G \end{gathered} \tag{2.7}$$

Furthermore, a property of valid probability distributions $\boldsymbol{p} \in \mathbb{P}$ states that

$$\boldsymbol{p} \in \mathbb{P} \Leftrightarrow \sum_{i=1}^{n} p_i = 1$$

Thus, for $\mathbb{P}$ to be convex, it must be shown that Equation 2.7 holds for $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{P}$, or:

$$\sum_{i=1}^{n} \left(\theta x_i + (1-\theta)\, y_i\right) = 1$$

$$\theta \cdot \left[\sum_{i=1}^{n} x_i\right] + (1-\theta) \cdot \left[\sum_{i=1}^{n} y_i\right] = 1$$

where, due to the constraint above, the sums are equal to one. Thus, the above equation simplifies to

$$\theta + (1-\theta) = 1$$

which is clearly true. Therefore, $\sum_{i=1}^{n} (\theta x_i + (1-\theta)\, y_i) \in \mathbb{P}$ and thus $\mathbb{P}$ is convex. □

Therefore, a proper scoring rule can use any convex cost function, G, that maps an entire distribution to a single real value. One such example is that of the *spherical* rule, whose cost function is

$$G(\mathbf{p}) = \left(\sum_{j=1}^{n} p_j^2\right)^{\frac{1}{2}} \tag{2.8}$$

and therefore

$$G'(\mathbf{p}) = \frac{\sum_{j=1}^{n} p_j}{\left(\sum_{j=1}^{m} p_j^2\right)^{\frac{1}{2}}} \tag{2.9}$$

From this, the spherical scoring rule can be derived by applying these functions to Equation 2.6, to get

$$\begin{aligned} S(\mathbf{p},\omega) &= \left(\sum_{i=1}^{n} p_i^2\right)^{\frac{1}{2}} - \mathbf{p} \cdot \left(\sum_{i=1}^{n} p_i\right) \cdot \left(\sum_{i=1}^{n} p_i^2\right)^{-\frac{1}{2}} + p_\omega \cdot \left(\sum_{i=1}^{n} p_i^2\right)^{-\frac{1}{2}} \\ &= \left(\sum_{i=1}^{n} p_i^2\right)^{\frac{1}{2}} - \left(\sum_{i=1}^{n} p_i^2\right) \cdot \left(\sum_{i=1}^{n} p_i^2\right)^{-\frac{1}{2}} + p_\omega \cdot \left(\sum_{i=1}^{n} p_i^2\right)^{-\frac{1}{2}} \\ &= \left(\sum_{i=1}^{n} p_i^2\right)^{\frac{1}{2}} - \left(\sum_{i=1}^{n} p_i^2\right)^{\frac{1}{2}} + p_\omega \cdot \left(\sum_{i=1}^{n} p_i^2\right)^{-\frac{1}{2}} \\ &= p_\omega \cdot \left(\sum_{i=1}^{n} p_i^2\right)^{-\frac{1}{2}} \end{aligned} \tag{2.10}$$

The derivation of the logarithmic and quadratic scoring rules from their respective cost functions is discussed in Appendix A. Further characterisations of these rules in terms of their properties have been provided by a number of authors, as is discussed next.

2.6.1.2 Properties of Discrete Scoring Rules

Although all of the above rules are strictly proper, they each provide differing incentives to agents. For example, in terms of fairness and rationality, the logarithmic scoring rule has no lower bound; the scores it awards agents are in the range $(-\infty, 0]$. This complicates matters when implementing a mechanism based upon scoring rules. As discussed earlier, agents must expect to profit through participation in a mechanism for it to be classed as *individually rational.* If there is no lower bound on the score awarded by the logarithmic scoring rule, ensuring positive payments to the players is non-trivial. On the other hand, the quadratic and spherical rules award scores with the bounds $[-1, 1]$ and $[0, 1]$ respectively (Bickel, 2007). Furthermore, the logarithmic score is said to exhibit *locality.* That is, the score rewarded by the logarithmic rule depends only on the precision allocated to the event that occurred. For example, take two reports forecasting the probability of three events, $\boldsymbol{a} = [0.7, 0.3, 0]$, and $\boldsymbol{b} = [0.7, 0.1, 0.2]$, if event 1 were to occur, both agents would receive the same logarithmic score, $\log(0.7)$, despite the fact they both reported two different distributions. The quadratic and spherical rules, on the other hand, exhibit a property known as *monotonicity*, whereby the agent receives a higher expected score for reporting a distribution that more closely matches the real distribution being predicted. These rules are said to be *effective* and encourage forecasters not only to report their beliefs truthfully but to generate forecasts that match the distribution of the event in question (Friedman, 1983). The range of scores awarded by the spherical rule when event ω has occurred and was predicted to occur with $\boldsymbol{p}(\omega)$ probability is much wider than the range of scores awarded by the quadratic rule under the same conditions (Bickel, 2007).

The rules thus far, although commonly used throughout the literature, make one limiting assumption – that no prior information is held by the person who requests the forecast. All reports are scored against a uniform baseline belief, and as such, actual skill and effort employed by the forecaster might not be taken into account by the score (Winkler, 1994). For example, consider the situation in which a commuter would like to get a prediction of whether or not a train will arrive within a given 10 minute period. Given the timetable and the train's history of being on time, the person knows there is a 70% chance of the train arriving within that 10 minute period. If the forecaster were simply to generate its prediction using this same information and therefore report the same probability to the commuter, that information will have no value to the commuter. However, if the forecaster were to gather other information, for example timing information from checkpoints along the train's route, and use that information to determine and report that there is in fact a 97% chance of the train arriving within that 10 minute period, this information has a much greater value to the commuter.

One scoring rule that has been used in meteorology to address this issue is the *skill score*, which is defined as:

$$S(\boldsymbol{p},\boldsymbol{q}|\omega) = \frac{T(\boldsymbol{p}|\omega) - T(\boldsymbol{q}|\omega)}{T(\{\omega\}|\omega) - T(\boldsymbol{q}|\omega)} \tag{2.11}$$

for a strictly proper scoring rule $T(\cdot)$, reported belief $\boldsymbol{p}$, event ω and, crucially, a base distribution $\boldsymbol{q}$. The term $T(\{\omega\}|\omega)$ refers to the score obtained when the forecast is exactly correct – i.e. the forecaster predicted a 100% chance of *rain* and the event $\omega = rain$. Now, it can be seen, that the skill score is the ratio of the difference between the score obtained using the reported belief $\boldsymbol{p}$ and the baseline distribution $\boldsymbol{q}$, to the difference between the score obtained when the reported belief is 100% accurate, and the baseline distribution. However, the skill score is *not* proper (Murphy, 1973), and therefore does not incentivise the forecaster to report their honest beliefs in all cases.

Work done in Winkler (1994) discusses the application of *asymmetric scoring rules*, which score a report against some base distribution. This work was extended by Nau et al. (2007) in order to develop a set of scoring rules named *weighted scoring rules*, which are based upon the generalised quadratic and spherical scoring rules. With these rules, given in equations 2.12 and 2.13, a forecaster's report $\boldsymbol{p}$ is scored against a realised event i and a base distribution $\boldsymbol{q}$.

$$\text{Weighted Spherical: } S(\boldsymbol{p},\omega|\boldsymbol{q}) = \frac{1}{\beta-1}\left(\left(\frac{\boldsymbol{p}(\omega)/\boldsymbol{q}(\omega)}{\mathbb{E}_{\boldsymbol{p}}\left((\boldsymbol{p}/\boldsymbol{q})^{\beta-1}\right)-1}\right)^{\beta-1} - 1\right) \tag{2.12}$$

$$\text{Weighted Quadratic: } S(\boldsymbol{p},\omega|\boldsymbol{q}) = \frac{(\boldsymbol{p}(\omega)/\boldsymbol{q}(\omega))^{\beta-1}-1}{\beta-1} - \frac{\mathbb{E}_{\boldsymbol{p}}\left((\boldsymbol{p}/\boldsymbol{q})^{\beta-1}\right)-1}{\beta} \tag{2.13}$$

With these rules, a forecaster's score is reduced not only in the case that he misreports, but also when he reports information already known by the centre. Consequently, the forecaster is incentivised to gather information to make $\boldsymbol{p}$ more precise than the information held by the centre. However, computing the weighted score is problematic and results in a score that is unbounded. For example, consider the situation in which a forecaster believes there is a 20% chance of *rain*, but the scorer believes there to be a 0% chance of *rain*, it can be seen that if $\omega = rain$, then $\boldsymbol{p}(\omega)/\boldsymbol{q}(\omega) = \infty$.

The early applications of scoring rules were mainly in domains with discrete outcomes. For example, weather forecasting where either it will rain or not rain. However, work has been done to adapt these rules for continuous domains (Winkler, 1967; Matheson and Winkler, 1976), as is discussed in the next section.

2.6.2 Scoring Rules for Continuous Distributions

As was discussed in the previous section, the preliminary work on scoring rules tended to study their application to discrete domains. However, recently much work has been done

Logarithmic	Quadratic	Spherical
$log\left(\mathcal{N}\left(\omega;\hat{x},\hat{\theta}\right)\right)$	$2\cdot\mathcal{N}\left(\omega;\hat{\mu},\hat{\theta}\right)-\int_{-\infty}^{\infty}\mathcal{N}\left(x;\hat{\mu},\hat{\theta}\right)^{2}\mathrm{d}x$	$\frac{\mathcal{N}\left(\omega;\hat{\mu},\hat{\theta}\right)}{\sqrt{\int_{-\infty}^{\infty}\mathcal{N}\left(x;\hat{\mu},\hat{\theta}\right)^{2}\mathrm{d}x}}$

Table 2.4: The logarithmic, quadratic and spherical scoring rules applied to a belief, $\hat{x}$, with mean, $\hat{\mu}$, and precision, $\hat{\theta}$, when the actual event, ω occurs.

regarding their application to continuous domains. This section discusses how discrete scoring rules can be adapted such that they can be used in continuous domains and the continuous variants of the logarithmic, quadratic and spherical rules discussed earlier are presented along with some other useful scoring rules.

Given a discrete proper scoring rule for a binary outcome $\mathcal{S}(\cdot)$, it was shown by Winkler (1967) that this rule can be used to score a forecaster in a continuous setting by setting a random threshold on the score obtained by the agent for a given outcome. For example, consider that:

$$\mathcal{S}(\hat{r})=\begin{cases}\mathcal{S}_1, & \text{if } \hat{r} \text{ occurs}\\ \mathcal{S}_2, & \text{otherwise}\end{cases}$$

By randomly partitioning the continuous outcome space, Ω, at a point u, proper scoring rules designed for binary domains can be adapted to a continuous scoring rule $S(\cdot)$, where

$$S(x|\omega)=\begin{cases}\mathcal{S}_1, & \text{if } \mathrm{sign}(\omega-u)=\mathrm{sign}(x-u)\\ \mathcal{S}_2, & \text{otherwise}\end{cases}$$

and where x is a point estimate provided by the forecaster and ω is the realised outcome (Winkler, 1967). However, in order to maintain incentive compatibility, it is imperative that u is unknown by the forecaster. To alleviate this problem, the work in Winkler (1967) was extended by Matheson and Winkler (1976) whereby u is integrated over Ω to produce a scoring rule that scores an agent based upon its reported continuous probability distribution with *cumulative* distribution function, r, over Ω, against an outcome $\omega \in \Omega$.

$$S(\hat{r}|\omega)=\int_{-\infty}^{\omega}\mathcal{S}_2(\hat{r}(u))\,\mathrm{d}u+\int_{\omega}^{\infty}\mathcal{S}_1(\hat{r}(u))\,\mathrm{d}u$$

This method relieves the partition selection problem in Winkler (1967), but also uses the entire distribution reported by the agent whereas previously only a point estimate was used. The discrete quadratic, logarithmic and spherical scoring rules described in Equations 2.2, 2.4 and 2.5 respectively, have also been adjusted to work in continuous domains. These are summarised in Table 2.4.

2.6.2.1 The Characterisation of Continuous Scoring Rules

The characterisation of continuous scoring rules, as provided by McCarthy (1956), Hendrickson and Buehler (1971), and restated in Gneiting and Raftery (2007), is very similar to that provided by Savage (1971) for discrete scoring rules. Again, it relies on the convexity of the scoring rule's associated cost function. Consequently, an adaptation of Theorem 2.1 stating the convexity of the space of discrete distributions, $\mathbb{P}$, is adapted for continuous distributions in Theorem 2.2.

Theorem 2.2. *The set of continuous probability distributions over an interval* $[a,b]$, $\mathbb{P}$ *is convex.*

Proof. The proof follows much the same lines as that of Theorem 2.1, and therefore will not be restated in full here. The proof for continuous distributions uses the property:

$$p \in \mathbb{P} \Leftrightarrow \int_a^b p\,(x)\,\mathrm{d}x = 1$$

and therefore

$$p, q \in \mathbb{P},\ \ \theta \in [0,1]$$

$$\theta \int_a^b p(x)\,\mathrm{d}x + (1-\theta) \int_a^b q(x)\,\mathrm{d}x = 1$$
$$\therefore (\theta p + (1-\theta) q) \in \mathbb{P}$$

□

Given the convex set of probability distributions, $\mathbb{P}$, the convex cost function $G : \mathbb{P} \to \mathbb{R}$, satisfies (R. Tyrrell Rockafellar, 1970):

$$\forall\, x, y \in \mathbb{P},\ \ \forall\, \theta \in [0,1] \tag{2.14}$$
$$G\,(\theta x + (1-\theta)\,y) \leq \theta G\,(x) + (1-\theta)\,G\,(y) \tag{2.15}$$

Furthermore, given the convex function, $G : \mathbb{P} \to \mathbb{R}$, the *subderivative* of G, is $G^* : \mathbb{P} \times \Omega \to \mathbb{R}$ such that:

$$G(y) - G(x) \geq \int_{\omega \in \Omega} G^*(x, \omega)\,\mathrm{d}\,(y-x)\,(\omega)$$

which can be derived from the standard form of a subderivative function (R. Tyrrell Rockafellar 1970, pp. 242;Gneiting and Raftery 2007). For a differentiable function G, the subderivative G^* can simply be the first derivative of G with respect to $\mathbb{P}$ – the use of *sub*derivatives simply allows for non-continuous cost functions.

A scoring rule, $S : \mathbb{P} \times \Omega \to \mathbb{R}$, is proper iff S is of the form:

$$S\,(p, \omega) = G\,(p) - \int G^*\,(p, \omega)\,\mathrm{d}p\,(\omega) + G^*\,(p, \omega)$$

Given this characterisation, the quadratic scoring rule can be derived using the cost function

$$G(p) = \int p(\omega)^2 \, \mathrm{d}\omega$$

taking the derivative with respect to p, gives

$$G^*(p) = \int 2p(\omega) \, \mathrm{d}\omega$$

thus, giving the score

$$S(p, \omega) = \int p(\omega')^2 \, \mathrm{d}\omega' - \int 2p(\omega')^2 \, \mathrm{d}\omega' + 2p(\omega) \tag{2.16}$$

$$S(p, \omega) = 2p(\omega) - \int p(\omega')^2 \, \mathrm{d}\omega' \tag{2.17}$$

Further derivations of the logarithmic and spherical scoring rules can be found in Appendix A.

2.6.2.2 The Continuous Ranked Probability Score

Matheson and Winkler (1976) and Gneiting and Raftery (2007) argue that the rules presented so far are not sensitive to *distance* – that is, they do not award high scores to forecasters that assign high probabilities to events close to the measured outcome but not identical to it. To this end, Matheson and Winkler (1976) propose a scoring rule named the *continuous ranked probability score* (CRPS), where given a reported distribution with *cumulative distribution function* F, and an outcome ω, the forecaster receives the score:

$$\mathrm{S}(F, \omega) = -\int_{-\infty}^{\infty} (F(x) - \mathbf{1}(x \geq \omega))^2 \, \mathrm{d}x \tag{2.18}$$

where

$$\mathbf{1}(predicate) = \begin{cases} 1, & \text{if } predicate \text{ is true} \\ 0, & \text{otherwise} \end{cases}$$

However, through the use of entire continuous distributions reported by the forecasters, many of the rules that have been discussed so far are in fact sensitive to distance insofar as the definition above describes. To reiterate, a scoring rule is said to be sensitive if it assigns high scores to forecasts that assign high probabilities to events close to the event that occurred (Matheson and Winkler, 1976; Staël von Holstein, 1970). It can be seen that in the discrete sense, with an outcome space of $\Omega = \{sun, rain, thunder\}$, a forecast that reports $\boldsymbol{r}_1 = \{0.7, 0.2, 0.1\}$ will receive the same score under the logarithmic rule

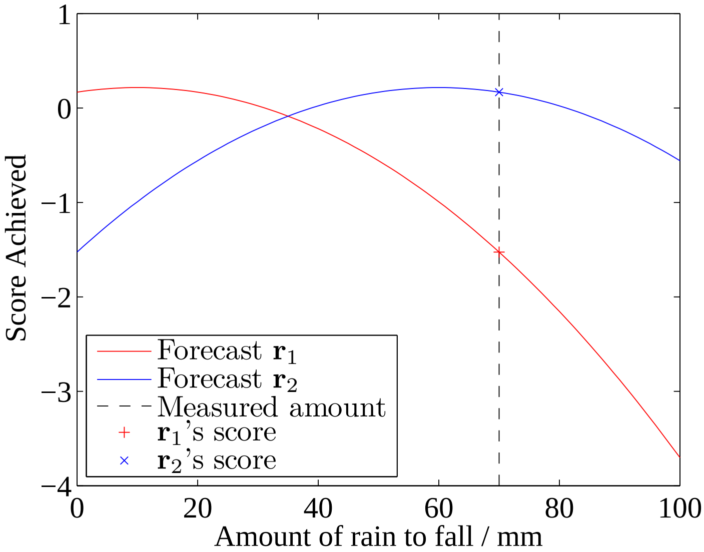

FIGURE 2.4: The logarithmic scores forecasters would receive for their predictions of how much rain will fall, $\boldsymbol{r}_1 = \mathcal{N}(0.1, 0.3)$ and $\boldsymbol{r}_2 = \mathcal{N}(0.6, 0.3)$. The dashed line shows the actual measured rainfall, and the intersection of that line with the curves is the score each forecast receives.

(for example) as a report that states $\boldsymbol{r}_2 = \{0.2, 0.7, 0.1\}$ when *thunder* occurs despite the fact that one report (giving a higher probability of rain) clearly is more accurate than the other. The logarithmic rule, therefore is clearly not sensitive to distance. However, in a continuous domain, when the forecaster is asked to give a Gaussian (for example) distribution of how much rain he believes will fall in a given day, it can be seen from Figure 2.4 that when two forecasts are given $\boldsymbol{r}_1 = \mathcal{N}(10, 0.3)$ and $\boldsymbol{r}_2 = \mathcal{N}(60, 0.3)$ and the actual measured amount of rain is 70mm, it is the 'closer' forecast that is more greatly awarded. Note that the logarithmic rule is still insensitive to distance, this is simply an outcome of using a distribution with a continuously differentiable probability density function.

It can also be seen in Figure 2.4 that forecasts receive a score greater than the upper bound of the logarithmic scoring rule that was specified in Section 2.6.1. As is the case with the quadratic and spherical rules, the translation of the logarithmic rule to the continuous domain destroys the boundedness of the rule. The problem arises due to the range of the continuous probability density function that describes the distribution being reported by the forecaster. There are methods that can be employed in order to scale the scores awarded by proper scoring rules. However, as is discussed in the next section, this is not always a trivial matter.

2.6.3 The Transformation of Scoring Rules

It is commonly accepted within the field of scoring rules that the affine transformation of proper scoring rules maintains their propriety. That is, given a strictly proper scoring rule $S(\cdot)$, the transformation $\alpha S(\cdot) + \beta$ is also strictly proper, for any positive linear α, and β. This is a property that has been routinely exploited by mechanism designers in order to develop scoring rule-based mechanisms that exhibit other properties as well as incentive compatibility. For example, in Miller et al. (2005) this technique is used to scale payments to the forecasters such that they will report a prediction with a given precision θ – an approach which is further developed by Papakonstantinou et al. (2008) from which the following explanation is adapted. In Miller et al. and Papakonstantinou et al.'s work, the agent's expected utility for reporting an estimate of precision, θ, at a cost, $c(\theta)$, is $\bar{U}(\theta) = \alpha\bar{S}(\theta) + \beta - c(\theta)$. In order to encourage agents to report an estimate of a given *required* precision θ_0, α must be carefully selected such that $\bar{U}$ is maximised when an agent reports $\theta = \theta_0$. This is done by solving $\mathrm{d}U/\mathrm{d}\theta|_{\theta_0} = 0$, resulting in:

$$\alpha = \frac{c'(\theta_0)}{\bar{S}(\theta_0)}$$

In order to satisfy Requirements II, III and IV – that agents are encouraged to participate, the aggregator does not lose money, and agents are incentivised to report their information truthfully – any payments made to the agents based upon scoring rules will likely have to be scaled similarly to above. However, the works in Miller et al. (2005) and Papakonstantinou et al. (2008) cannot be directly applied to the problem in this thesis as they make a limiting assumption that the payments to the forecasting agents are independent of their reports, and that multiple agents are reporting information on a common event.

It should be noted that arbitrary transformations are not guaranteed to maintain incentive compatibility. Any transformation of a scoring rule in effect creates a new scoring rule, which must still obey the characterisation discussed in Section 2.6.2.1 in order to be proper. However, even strictly proper scoring rules only provide guarantees regarding the *truthfulness* of agents' *reports*. Agents may well be able to strategise over other parameters, resulting in behaviours that, while still truthful, are detrimental to the system. A recent body of work has focussed on designing scoring rules that discourage these behaviours, as is discussed next.

2.6.4 Scoring Rules That Discourage Undesirable Actions

A common assumption among the scoring rule and information elicitation literature is that an agent is unable to influence the event which it is predicting. However, if this assumption is relaxed, as it is in the problem described in Chapter 3, interesting new incentives are introduced. In more detail, an agent who is able to predict the value of

ω	$\mathbf{p}(\omega)$	$S^L(\mathbf{p},\omega)$	Bonus
October	0.2	-0.699	\$1.
November	0.8	-0.0969	\$3.

TABLE 2.5: Scores and bonuses (taken as $\left[10 \cdot S^L(\mathbf{p},\omega)+10\right] \cdot c$) awarded to a software developer who predicts a $\mathbf{p}(\omega)$ chance of software being ready to ship in $\omega \in \{October, November\}$.

an event and knows its reward will be maximised in the case that it is correct may have incentive to manipulate its event such that it matches the agent's report, if the agent were to discover that its initial estimate was false. This is more concretely described in the following example.

Consider, for example, a software firm that would like to determine when the next release of their product will be ready for shipping. To do this, they request a prediction from their software developers as to whether the product will be available in $\omega \in \Omega = \{October, November\}$. For simplicity in this example, let the software developers receive a bonus based on some affine transformation of the logarithmic scoring rule – taken in this case to be $\left(10 \cdot S^L(\mathbf{p},\omega)+10\right) \cdot c$, for some constant, c. An example of the bonuses the software developer receives on predicting a 0.2 probability of the software being ready in October is shown in Table 2.5. It can be seen, unsurprisingly, that the bonus received by the developer is maximised if his prediction is correct. Now, imagine that after the developer has reported his belief, he finds that software development is progressing much faster than expected, and he is almost certain that, if progress continues at the current rate, the software will be ready in October. Furthermore, he knows that if this is the case, he will receive only \$1 instead of \$3 as his bonus. Therefore, he stalls development of his part of the software such that it is guaranteed to be complete in November, thereby increasing his expected reward.

Clearly, in this situation, this type of behaviour from the agent is undesirable. However, it should be noted, that under some circumstances this is not the case. For example, in the event that the developer discovers that he is working too slowly and so speeds up his development in order to ensure the software is delivered at the time that he estimated.

Scoring rules that discourage this type of behaviour are termed *principally aligned* (Shi et al., 2009). That is, the scores assigned by the scoring rule are aligned with the utility function of the principal agent. Thus for the agent to maximise its own score, it must act to maximise the utility of the principal. In their paper, Shi et al. (2009) provide a characterisation of principal aligned scoring rules, and discusses their use in *prediction markets*.

Further work along these lines was presented by Bacon et al. (2012) wherein a model was developed in which a worker must send information to a manager about the expected time it will take to complete a given task. The manager has prior information, which it fuses with the information from the worker in order to make a prediction as to when a task will be completed. This prediction is then submitted to the company, who will then reward both the manager and the worker proportionally according to a set of scoring rules S_m and S_w respectively. It is assumed that there is a maximum amount of effort that can be attained by the worker, and in exerting this effort, the task will be completed in some time, x. By definition, a worker cannot exert more than maximum effort, and therefore x represents the shortest amount of time in which the task can be completed.

The authors provide a characterisation of a set of scoring rules that incentivise the worker to work at best effort, and truthfully report its information regarding the estimated completion time to its manager. The scoring rules further provide incentives for the manager to use all information available to it (including the information reported by the worker) in order to generate its own estimate as to when a project will be completed, and then to truthfully report that mean time to the company.

However, this assumes no cost of effort on the part of the worker, who, in absence of a reward mechanism, is indifferent between exerting minimal and best effort. Moreover, no thought is given to the rewards that are generated under such a mechanism – just that they are distributed in such a way that agents are incentivised to maximise their score. Splitting the reward between the worker and the agents proportionally based upon their score is likely to be non-trivial as the worker is able to manipulate the manager's estimate through manipulation of the information sent by the worker to the manager.

An interesting situation that might arise in the scenarios discussed in this thesis occurs when agents speculate as to the action that is to be chosen by the centre, and how that action will affect the agent's utility. This problem is effectively side-stepped by assuming that agents' utilities are independent of the actual action chosen by the centre, provided the centre acts optimally according to the information it holds. That is, using the electricity domain as an example, a customer is indifferent between the centre purchasing ten units or twenty units of electricity provided whichever amount the centre chooses to buy minimises the total cost of electricity to the customer. However, it has been shown that when agents receive additional utility dependent on the actual action chosen by the centre – e.g. when the customer in the previous example receives some sort of commission from the provider from which the centre buys its electricity – the centre must provide additional reward – compensation – to the home agent in order to ensure truthful reporting of beliefs (Boutilier, 2012).

2.7 Aggregation of Loads and Coalition Formation

Recently, work done in Chalkiadakis et al. (2011) employed coalition formation in order to effectively aggregate small energy generators, such as wind farms or solar arrays, known as distributed energy resources (DERs), in order to form a collective known as a virtual power plant. The coalition then acts as a whole to report to the grid the amount of electricity it expects to generate in the future, and then to generate and provide that electricity to the grid when the time arises. In return, after the time in question has passed, the coalition is paid an amount by the grid that:

$$V_{G,C} = \frac{1}{1 + \alpha \left| \hat{x}_C - x_C \right|^{\beta}} \cdot \log \left(x_C \right) \cdot \kappa \cdot x_C$$

where $\hat{x}_C$ is the amount of electricity the coalition claimed it would generate, x_C is the *actual* amount generated by the coalition, and κ is the price per unit of electricity as set by the grid operator. They show that for the coalition, this method of payment is incentive compatible. That is, as a whole, the coalition's report of $\hat{x}_C$ should be equal to x_C. However, the payments made to individuals within the coalition must be made carefully such that the mechanism is *individually* incentive compatible. The payment mechanism they use to pay each agent i within the coalition is given as follows:

$$V_{C,i} = \frac{z}{1 + \alpha \left| \hat{x}_i - x_i \right|^{\beta}} \cdot \frac{x_i}{x_C} \cdot V_{G,C}$$

where $\hat{x}_i$ is the amount of electricity agent i claimed it would generate, and x_i is the actual amount it generated. This mechanism is individually incentive compatible. That is, individually, each agent will maximise its reward if it reports its expected generation capacity truthfully. The cited work, however, does not take into account the *variance* of the agent's belief that it will produce $\hat{x}_i$ units of electricity, and as such, only works in situations whereby the cost of error is symmetric (which generally is not the case in electricity markets). Further work, in Robu et al. (2012), which was performed in parallel with the work in this thesis, builds upon Chalkiadakis et al. (2011) by looking at a situation in which penalties are asymmetric, and therefore the variance of each agents' belief is actually utilised by the coalition. In this work, the continuous ranked probability score (given in Equation 2.18) is used in order to determine the amount to be given to the coalition, and again to determine the amount given to each individual within the coalition. In this adapted mechanism, the amount paid to the coalition by the grid for the energy produced by that coalition is given by:

$$V_{G,C} = \frac{\text{CRPS}\left(\mathcal{N}\left(\mu_C, \sigma_C^2\right), x_C\right) + x_C}{x_C} \cdot \log\left(x_C\right) \cdot \kappa \cdot x_C \tag{2.19}$$

where μ_C and σ_C^2 are respectively the mean and variance of the distribution reported by the coalition describing how much electricity it estimated it would be able to generate. The payments to individuals within the coalition then becomes:

$$V_{C,i} = \frac{\text{CRPS}\left(\mathcal{N}\left(\mu_i, \sigma_i^2\right), x_i\right) + x_i}{\Sigma_{\forall j \in C}\left(\text{CRPS}\left(\mathcal{N}\left(\mu_j, \sigma_j^2\right), x_j\right) + x_j\right)} \cdot V_{G,C}$$

However, a disadvantage to this system is that it requires a redesign of the electricity markets currently in use (and therefore breaks Requirement XI). The equation in Equation 2.19 shows the amount that the grid would pay each the coalition for the electricity they supply. However, this does not reflect the current payment system in use, as discussed in Section 2.1. A mechanism that is more immediately applicable to the electricity market should be able to fit into the current system, without redesigning the market, which, without thorough testing can have costly consequences (Borenstein, 2002).

2.8 Summary

To summarise, this chapter has presented the required background to understand the remainder of this thesis. In addition, the current state of the art in terms of mechanism design and information elicitation has been discussed. It was shown that while much work has been done on eliciting information from agents, certain limitations reduce its applicability to the scenario presented in this thesis. Namely, these limitations are:

- Work thus far has not taken into account forms of manipulation possible within the model from Chapter 3. In particular, no consideration has been made to providing incentives for agents who have the ability to make themselves harder to predict.
- Mechanisms presented so far in literature do not take into account agents' incentives to fix their events when they have discovered they have reported a false prediction *and* agents incur a cost of effort. Whereas Bacon et al. (2012) provides incentives for agents not to fix their outcomes, they make the assumption that an agent is indifferent between working at best effort in order to achieve their outcome, or working at some other rate, less than best effort. In the demand prediction problem, the home agents incur a cost in consuming electricity, thus it is not clear that the rules in Bacon et al. (2012) are applicable to the demand prediction scenario.
- Furthermore, work which makes use of scoring rules in order to allocate rewards does not take into account budget balance. That is, payments towards agents based upon their scores may result in the centre making a loss. In some cases, payments themselves are not discussed at all, but instead are assumed to be allocated in such a way that agents are incentivised to maximise their score. The work in this thesis allocates payments to agents based upon the agents' scores in a way that is *budget balanced*, meaning the centre does not make a loss through payments.
- Mechanisms so far tend not to take into account prior information held by the centre. Consequently, the centre may pay agents for information it already holds.

The mechanisms in this thesis are designed such that the centre only pays agents who report more information than is currently held by the centre.

- Many information elicitation mechanisms rely on the fact that there are a number of agents from whom the centre is able to request information. They then make use of this fact by scoring agents by comparing their reports to one-another, or fusing information in order to get a precise report. In the demand prediction problem, it is assumed that the overall outcome is dependent on a set of events, information of which is private to each agent. Therefore, only a single agent is able to provide information regarding each event and no agents are able to provide information on the overall outcome.

- Mechanisms for information elicitation tend to use fixed budgets. However, the mechanisms in this thesis use a budget that is itself dependent upon the agents' reports. This makes the allocation of rewards even less trivial, as agents may be able to either collude with one-another or manipulate their consumptions or reports in order to gain a greater payment.

Given these limitations and the background presented within this chapter, the remainder of this thesis is dedicated to the presentation of the model of the problem on which this thesis focuses and the design and analysis of the mechanisms that achieve the properties described in Section 1.1.

Chapter 3

Information Aggregation in the Smart Grid

This chapter discusses the model of the information aggregation problem studied in this thesis. The basic model is that of the *newsvendor problem* discussed in Section 2.3.2, extended with the ability for the newsvendor to gather information from its clients regarding their future behaviours. This chapter begins with a scenario in which an electricity company must purchase electricity for a consumer. The aim is to provide the reader with a concrete example on which to ground the more abstract discussions later in this thesis, such that concepts may more easily be understood. Following the scenario, Section 3.2 provides a formal model of the information aggregation problem. Then, Section 3.3 discusses ways in which agents my attempt to game the system through manipulation in order to gain higher rewards. Section 3.4 discusses the home agents' types and strategies, followed by a discussion of theirs and the aggregator's utility functions in Section 3.5. Finally, Section 3.6 discusses computational aspects of the model, which must be taken into account during empirical evaluation of the solutions presented in later chapters.

3.1 Scenario

Alice has recently moved house and has joined a new electricity aggregator who will buy electricity from the grid and supply it to her, charging a fixed rate of 13p per kWh. The electricity company buys a day's worth of electricity at a time for Alice. In the day ahead, it buys the electricity from the 'forward market', whose prices are set at 7p per kWh to buy and sell electricity. Should Alice consume more energy than was initially bought for her from the grid, the electricity company must buy more electricity to top up the difference. As the aggregator now has to buy this electricity in 'real-time', it must do so from a balancing market. To encourage better planning of electricity purchasing, the balancing *market's rates are set by the grid operator to 10p for one kWh of electricity. Likewise, should Alice consume less energy than was bought for her, any unused electricity must be sold back to the grid. This energy must also be sold through the balancing market, and it is done so at the rate of only 4p per kWh.*

In order to minimise the cost of electricity, the company must buy exactly the required electricity a day in advance. The electricity company can reduce the net error of the amount of energy they predict to be consumed by simply purchasing energy for more houses. This strategy reduces the total error in the amount of electricity that will be consumed due to numerous errors in individuals' predictions cancelling each other out. For example take Alice, who is predicted to use 10kWh; Bob, who is predicted to use 5kWh; and Claire, who is also predicted to use 10kWh. All 25 kWh of electricity has been bought by the electricity company 24 hours in advance, but the reported predictions were not accurate. On the day, Alice used 7kWh, Bob used 10kWh and Claire used 5kWh. Had they bought their electricity individually, Alice would have paid 8.2p per kWh, Bob 8.5p per kWh and Claire 10p per kWh. By aggregating, the company buys 25kWh, but only uses 22 kWh. It can be clearly seen in this example that Bob and Claire's net error is zero, meaning the overall error is only 3 kWh resulting in a net cost of 7.4p per kWh.

The aggregation service can gather consumption information about the homes for which it is responsible, which it can use to make predictions about future energy consumption. The fact that Alice has only just joined this aggregation service means that they have no data on her past consumption, and so they are not able to accurately predict the amount of electricity she will use. Alice, however, knows her own comfort preferences and schedule, and has been provided with a smart meter to which she can supply this information. By using this information, as well as information about the devices in her property, software running on Alice's smart meter is able to predict to within a reduced margin of error the amount of energy she will use 24 hours in advance. For example, it knows that tomorrow Alice has a day off work, which she will be spending at home. It knows that, as it is the middle of winter and Alice likes to keep her home at around 21° C, her electric heating is likely to consume a lot of power compared to when she is out of her house. The smart meter can then tell the aggregator how much energy Alice will use and if it is accurate (i.e. after the time for which she has predicted her consumption, her real consumption closely matches her prediction), Alice will be rewarded by being charged a reduced cost of electricity per unit by the aggregator.

One month later, Alice's electricity provider is better able to predict the amount of electricity she will consume in a day by comparing against her previous consumption. If Alice's meter simply reports the amount of electricity she uses based on what she has used before, then the electricity company will not gain anything by using Alice's information over their own. However, Alice will always have more specific information about her consumption than the aggregator does. This information exists in many forms such as computerised calendars, comfort profiles and house use. By supplying this information *to her smart meter, she can still benefit from a reduced electricity cost. For example, if from tomorrow Alice will be on holiday for a week, her consumption for this period will be minimal - 2kWh for a day. The aggregator is expecting Alice to use 10kWh, and so knowing that Alice will only use 2kWh rather than 10kWh will save the aggregator money as they won't have to sell the excess 8kWh of electricity back to the grid at a loss.*

Once Alice returns from her holiday, she invites her friends over for a party on the weekend. Her meter then knows that she is going to use a lot of electricity, but it does not know the exact figure that she will consume. Therefore when it reports Alice's expected consumption to the electricity company, it reports that there is a 25% chance she will consume 15–18kWh of electricity, 50% chance that she will consume 12–15kWh, a 20% chance that she will use 10–12 kWh, and a 5% chance that she will use less than 10kWh of electricity. The aggregator encourages consumers to report their uncertainty and uses this information in calculating their rewards after the event. In fact, the consumers must report their uncertainty in order to be rewarded by the mechanism. After receiving the above distribution of Alice's meter's beliefs, the aggregator does some processing and determines that the expected amount of electricity Alice will consume is 14.9kWh, which it then buys for her. When the amount Alice consumed during that period is actually measured, the total comes to only 12kWh, and therefore Alice does not receive as great a reward as she would have done if she had used exactly 14.9kWh. The reward Alice receives is carefully calculated to discourage her from upwardly adjusting her consumption to match an earlier prediction which proved to be too high. However, Alice has noticed that as time progresses the reward she receives is reduced because the aggregator is learning more about her consumption and as such her reports are becoming less useful to it. To avoid this, she tries to make her consumption harder to predict such that her information is more valuable to the aggregator.

3.2 Formalised Model

Given the scenario in the previous section, let there be a single aggregator agent, a, whose job it is to purchase electricity for n homes, each represented by an individual home agent. Let the set of all home agents be $N = \{1, \cdots, n\}$. The aggregator must choose an amount of electricity to buy, χ, for a future time such that the expected cost of that electricity is minimised. To do this, the aggregator can make use of the *forward market*, in which electricity is bought and sold at a price of f per unit. However, the centre is only able to use this market *ahead of time*. Any imbalance between the amount of electricity the consumers use and the amount of electricity bought for them by the aggregator must be rectified in real time in the *balancing market*. The costs in the balancing markets penalise the aggregator for imbalance by forcing the aggregator to purchase extra electricity at the *system sell price*, $f + \delta^s$, or to sell back excess electricity at the *system buy price*, $f - \delta^b$.

Each home agent, i, represents a home, and requires to place some stochastic load, $\ddot{\omega}_i$ upon the grid. However, agents are able to strategise over the amount of electricity they consume, and therefore the load actually placed on the grid by that home is ω_i. Let the vector of the agents' consumptions be $\boldsymbol{\omega} = \{\omega_1, \cdots, \omega_n\}$ and let the sum of all agents' consumptions in $\boldsymbol{\omega}$ be ω.

Given this, if the aggregator purchases χ units of electricity from the forward market, and the consumers use in total ω units of electricity, the total cost of the electricity is:

$$\kappa(\chi, \omega) = f \cdot \chi + (\omega - \chi) \cdot \begin{cases} (f + \delta^s) & \text{if } \omega > \chi \\ (f - \delta^b) & \text{otherwise} \end{cases} \tag{3.1}$$

In order to determine the optimal amount of electricity to purchase for the home agents, the aggregator must make a prediction of the agents' future consumptions. It is able to do this based upon historical records of each agent's electricity use. Let $\Pi(\theta_{a,i})$ be some generic probabilistic distribution with information $\theta_{a,i}$, representing the amount of information held by the aggregator about agent i's future consumption. The aggregator therefore has a set of beliefs about each agent's future consumptions:

$$\boldsymbol{x}_a = \{x_{a,1}, \cdots, x_{a,n}\}$$

where

$$x_{a,i} \sim \Pi(\theta_{a,i})$$

Given these beliefs, the aggregator is able to make a prediction about the total consumption of the agents. Let x_a^+ be the aggregator's prediction of the total consumption of the home agents, with information, θ_a^+. Furthermore, let $\Pi(\theta_a^+)$ be some distribution, which is *not necessarily* the same as for individual agents[1]. With this, the aggregator is able to calculate the optimal amount of electricity to buy for the home agents (i.e. the amount of electricity to purchase that minimises the aggregator's expected cost):

$$\chi(\boldsymbol{x}_a) = \arg\min_x \int \kappa(x, \omega) \cdot \Pi\left(\omega; \theta_a^+\right) \, \mathrm{d}\omega \tag{3.2}$$

where $\Pi(\omega; \theta_a^+)$ is the probability density at ω of a generic distribution with information, θ_a^+.

For the sake of further discussion, let us assume a Gaussian distribution with mean μ^+ and precision θ^+, where precision is one over the variance of the distribution. In this way, the precision is analogous to the *Fisher information* of the distribution (Savage, 1976). Now, let the expected cost given some distribution, $x \sim \mathcal{N}(\mu^+, \theta^+)$ and an amount to

[1]The distribution of the total consumption of all agents is in fact the convolution of the probability density functions of the beliefs of the individual agents' future consumptions.

buy in the forward market, χ, be

$$\begin{aligned}\bar{\kappa}(\chi, x) = f \cdot \chi - &\int_{-\infty}^{\chi} \mathcal{N}\left(\omega; \mu^{+}, \theta^{+}\right)(\chi - \omega) \cdot \left(f - \delta^{b}\right) \, \mathrm{d}\omega \\ + &\int_{\chi}^{\infty} \mathcal{N}\left(\omega; \mu^{+}, \theta^{+}\right)(\omega - \chi) \cdot \left(f + \delta^{s}\right) \, \mathrm{d}\omega\end{aligned} \tag{3.3}$$

It can be seen from Equation 3.3, that the expected cost of electricity can be reduced by increasing the amount of information contained within the distribution. However, the aggregator only has access to a limited amount of information based upon each home's historical consumption. Conversely, each house has a wealth of information within it that can be used to make predictions. Consequently, if the aggregator were to use the home's information, it would be able to make some savings in the cost of electricity.

The aggregator is able to obtain this information by requesting predictions from each home agent about that individual home's future consumption. Each home agent is then able to generate a prediction, $x_i \sim \Pi(\theta_i)$, where θ_i represents the amount of information held by the agent. Home agents have access to their historical information as well as other, more detailed information within their respective homes. Consequently, it can be assumed that $\theta_i \geq \theta_{a,i}$. However, while maintaining historical information is assumed to be free, gathering more specific information is costly for the home agents. With this in mind, it is reasonable to assume that agent's costs are in some way proportional to the amount of *extra* information they provide over their historical estimates. That is, given an agent's information cost coefficient, $\alpha_{\theta,i}$, which represents the differing costs of effort exerted by the agents, an agent incurs the cost:

$$C_{\theta}\left(\theta_i, \alpha_{\theta,i}\right) = \alpha_{\theta,i} \cdot \left(\theta_i - \theta_{a,i}\right), \quad \theta_i > \theta_{a,i} \tag{3.4}$$

These costs can arise from a number of sources. For example, the cost of precision might be incurred through the cost of purchasing sensors to detect environmental variables within the property. Alternatively, the costs could arise from powering the extra sensory equipment required by the home agent. However, these costs are likely to be relatively small compared to the cost per unit of electricity, as the hardware costs will be constant with respect to the number of units of electricity the home consumes. Therefore, these costs will be divided across all units of electricity consumed by the house. Moreover, the cost of powering such sensor nodes will likely be small in comparison to the cost of electricity, as the hardware should be designed to consume as little power as possible.

When reporting their predictions to the aggregator, as is discussed in more detail in Section 3.3.1, each home agent reports a belief, $\hat{x}_i$, which is not necessarily equal to x_i. Crucially, this means that an agent's reported belief may either have a false expected value, or may report to contain more information, $\hat{\theta}_i$, than it truly does.

The aggregator builds an aggregate belief vector, $\boldsymbol{x} = \{\hat{x}_1, \cdots, \hat{x}_n\}$ and then uses the convolution of these beliefs in order to determine the amount to buy, as in Equation 3.2. As with the aggregated priors, let $\hat{\theta}^+$ represent the total information held by the aggregate belief vector. In doing this, the aggregator makes some savings compared to the amount it would have spent if it had used only $\boldsymbol{x}_a$, which is calculated as follows:

$$\Delta\left(\chi, \chi_a, \omega\right) = \kappa\left(\chi_a, \omega\right) - \kappa\left(\chi, \omega\right) \tag{3.5}$$

where

$$\chi = \arg\min_x \int \kappa\left(x, \omega\right) \cdot \Pi\left(\omega; \theta^+\right) \, \mathrm{d}\omega \tag{3.6}$$

$$\chi_a = \arg\min_x \int \kappa\left(x, \omega\right) \cdot \Pi\left(\omega; \theta_a^+\right) \, \mathrm{d}\omega \tag{3.7}$$

For their information, the agents are paid some reward by the aggregator. This is in part to encourage them to make predictions, as the agents are assumed to be rational, and in addition, to encourage desirable behaviour from the agents. That is, the agents receive a payment, $P : \mathcal{P}\left(\mathbb{P}^n\right) \times \mathbb{P}^n \times \Omega \rightarrow \mathbb{R}$. However, if this reward is not carefully designed, home agents may be able to manipulate either their reports, or their consumptions in order to gain extra money from the aggregator.

3.3 Manipulation within the Model

While the previous section discussed the basic model that underlies the problem focused upon within this thesis, the real problem to be solved is how to design payments such that the home agents behave in a manner beneficial to the aggregator. In this sense, naively designed payments may encourage agents to *misreport* their beliefs, increase the volatility of their loads, or even waste electricity. This section discusses these methods of manipulation in detail, beginning first with how agents may misreport their beliefs.

3.3.1 Misreporting

The simplest form of manipulation, and the one most commonly studied among mechanism design literature is to *misreport* the information held by the agent. Any payment mechanism that ensures agents report their information truthfully (their *type*, to use terminology from mechanism design) is said to be *incentive compatible*.

Careful design is required to design incentive compatible mechanisms as agents are able to misreport their information at *no extra cost*. Consider, for example, a naive reward mechanism, which pays agents for the amount of information they report. That is,

$$P(\boldsymbol{x}, \boldsymbol{x}_a, \omega) = \hat{\theta}_i.$$

Clearly, under such a mechanism, an agent can claim its report contains more information in order to get a larger payment. Moreover, as the realised outcome is not taken into account, such a payment mechanism makes agents indifferent between reporting the *expected value* of their distribution truthfully or misreporting it.

However, even mechanisms that take into account the realised outcome of an event are not guaranteed to be incentive compatible. Take, for example, the mechanism that pays agents based upon the squared error between their report and their realised outcome:

$$P(\boldsymbol{x}, \boldsymbol{x}_a, \boldsymbol{\omega}) = k - (\omega_i - \mathbb{E}(\hat{x}_i))^2$$

where $\mathbb{E}(\hat{x}_i)$ represents the expected value of the distribution $\hat{x}_i$. This mechanism encourages truthful reporting of the expected value of the agent's information. However, the agent is now indifferent regarding the truthful reporting of θ_i.

3.3.2 Load Variance Manipulation

Further to misreporting their information, agents may also attempt to game the system by making themselves harder to predict. They achieve this by making their loads more variable over time, and thus this type of manipulation is termed load variance manipulation. Let the variance of an agent's load in *the absence* of manipulation be termed it's *natural load variance*, or $\sigma^2_{nat,i}$. An agent is able to strategise over its *load variance manipulation factor*, γ_i, which affects the information contained within the aggregator's priors as follows:

$$\theta_{a,i} = \frac{1}{\sigma^2_{nat,i}} \cdot \begin{cases} (1+\gamma_i)^{-1} & \text{if } \gamma_i > 0 \\ (1-\gamma_i) & \text{otherwise.} \end{cases} \tag{3.8}$$

Note that load variance manipulation *does not* constitute misreporting. The agents physically manipulate their loads in order to make themselves harder to predict. Thus, applying load variance manipulation directly affects the amount of information held in the aggregator's priors. Moreover, this gives rise to relatively high costs (in comparison to the cost of gathering information of a given precision) for the agents. This cost may, for example, arise from the inconvenience of the home-owners having to modify their behaviour to achieve the more, or less volatile load. Each agent, i, has a load variance manipulation cost coefficient, $\alpha_{\gamma,i}$, and incurs the cost:

$$C_\gamma(\gamma_i, \alpha_{\gamma,i}) = \alpha_{\gamma,i} \cdot \gamma_i. \tag{3.9}$$

Mechanisms that pay agents based upon the additional information they provide, including the mechanisms discussed later in this thesis, must pay careful attention to this form of manipulation. Agents may be able to increase their reward by making themselves harder to predict, which in turn increases the difference in the amount of information they provide compared to the historical information held by the aggregator. Although it should be noted that, given Equations 3.4 and 3.8, an agent making itself harder to predict also increases the amount of effort it requires in order to gather its own precise information. Nevertheless, if the relative cost of gathering information is small, agents may still have an incentive to manipulate.

Given that agents are able to control their loads, it is reasonable to assume that they will not only manipulate the variance of their loads, but they will also actively fix the amount they consume in order to meet their predictions, as is discussed next.

3.3.3 Load Fixing

Most often, in mechanism design it is assumed that an agent's type is fixed or that there is a single period in which an agent can learn its type. However, in the case of the home agents presented here, agents are assumed to be able to continually gather information regarding their future consumptions, even once they have reported their information to the aggregator. Nonetheless, there is only a single opportunity for the agents to report the information they have to the aggregator. Consequently, an agent might learn that it is not going to consume as it had originally predicted, and will therefore, if the mechanism specifies, incur some penalty. For example, consider a mechanism that rewards agents based upon the score they receive:

$$P\left(\boldsymbol{x}, \boldsymbol{x}_a, \boldsymbol{\omega}\right) = S\left(x_i, \omega_i\right).$$

An agent will maximise its score if it consumes the expected value of x_i. Therefore, if it learns that $\omega_i < \mathbb{E}\left(x_i\right)$, it is potentially in the agent's interest to waste electricity such that $\omega_i = \mathbb{E}\left(x_i\right)$. Of course, in reality, the agent's updated belief given new information will itself be a distribution over Ω rather than the point estimate, ω_i. However, for the purpose of analysis, the situation in which an agent provides a point estimate represents that of an agent being certain of its future consumption, and therefore is useful for observing the worst-case behaviour of agents.

It is assumed that when wasting electricity, the only cost incurred by agents is that of the extra electricity consumed. This seems a reasonable assumption as home-owners are unlikely to be inconvenienced by an appliance running when it would otherwise not.

Load fixing is not always undesirable. Consider the situation in which an agent learns it will *over-consume*. Assuming a green agenda on behalf of the aggregator (or at least indifference over the amount of electricity the homes consume), it is desirable for agents to shed loads upon learning that they will consume more than predicted. In this sense, mechanisms designed for the smart grid should encourage the efficient use of electricity; penalising over-consumption but not under-consumption on the part of the home agents.

3.4 The Home Agents' Strategies and Types

In terms of agents' strategies, agents may choose to vary the amount of effort they place in gathering information, which in turn varies the informational content of their reports, $\theta_i \in \mathbb{R}^+$. In addition, agents may choose to misreport their information, thereby reporting to the centre some value, $\hat{x}_i = G_i(x_i)$, where $G_i : \mathbb{P} \to \mathbb{P}$, and $G \in \mathbb{G}$. Moreover, agents are able to make themselves harder or easier to predict by changing their load variance manipulation factor, $\gamma_i \in \mathbb{R}$, and are able to directly manipulate the amount of electricity they consume, $\omega_i \in \Omega$, but *not* the amount of electricity they *require*, $\ddot{\omega}_i$. Thus, the set of strategies for agents is given by:

$$\Sigma = \mathbb{R}^+ \times \mathbb{G} \times \mathbb{R} \times \Omega \tag{3.10}$$

The agents' types are dependent on the amount of electricity they *require*, $\ddot{\omega}_i \in \Omega$ the costs they incur for gathering information, $\alpha_{\theta,i} \in \mathbb{R}^+$ and manipulating their events, $\alpha_{\gamma,i} \in \mathbb{R}^+$. That is:

$$T = \Omega \times \mathbb{R}^+ \times \mathbb{R}^+$$

It is important to note that in terms of incentive compatibility, load variance manipulation and load fixing *are not* examples of misreporting. Thus, although a mechanism might be incentive compatible, agents may still perform these types of manipulation in order to gain extra utility. The agents' behaviours are dependent on the payment functions implemented by the mechanism as well as the individual agents' utility functions, which are discussed next.

3.5 The Home Agents' and Aggregator's Utility Functions

Given the underlying model described in Section 3.2 and the forms of manipulation discussed in Section 3.3, this section now presents the utility functions of the aggregator and home agents within the model.

3.5.1 The Aggregator's Utility Function

The aggregator gains utility from selling electricity to the home agents, which it does at a cost (known as the retail price) of f_r per unit consumed. Conversely, the aggregator loses utility by rewarding the agents for their information and purchasing the required amount of electricity from the markets. Formally, the aggregator's utility function is:

$$U_a\left(\boldsymbol{x}, \boldsymbol{x}_a, \boldsymbol{\omega}\right) = \left[\sum_{i=1}^{n} \left(f_r \cdot \omega_i - P_i\left(\boldsymbol{x}, \boldsymbol{x}_a, \boldsymbol{\omega}\right)\right)\right] - \kappa\left(\chi\left(\boldsymbol{x}\right), \omega\right) \tag{3.11}$$

3.5.2 The Home Agents' Utility Functions

The home agents' utility functions are more complex due to their ability to strategise. An agent gains utility from its reward and consuming electricity up to the amount required by the residents of the home, but it loses utility through payment for the electricity it consumes, effort exerted in gathering information and effort exerted in manipulating the volatility of its loads.

In addition, in the case of load fixing, it is necessary to model the utility gained by an agent in consuming each unit of electricity. If this was not taken into account in the agents' utility function, when performing analysis, home agents would simply choose to always consume zero units of electricity, as this minimises their costs. Therefore, to take this into account, it is assumed that each agent, i, requires a certain amount of electricity, $\ddot{\omega}_i$, and the utility it gains for consuming ω_i is given by:

$$P_\omega\left(\omega_i, \ddot{\omega}_i\right) = \begin{cases} f_r \cdot \omega_i & \text{if } \omega_i \leq \ddot{\omega}_i \\ f_r \cdot \ddot{\omega}_i & \text{otherwise} \end{cases} \tag{3.12}$$

That is, while an agent is consuming less than or equal to the amount it requires, it gains utility that is equivalent to that electricity's cost. However, once an agent consumes the amount it requires, it gains no extra utility in consuming more. This seems a logical assumption to make. In essence, the required electricity is the load that the residents of the home place on the grid. Equation 3.12 makes the assumption that the humans believe placing their load upon the grid is worth the cost of that load, and therefore that $f_r \cdot \ddot{\omega}_i = f_r \cdot \omega_i$. However, if the agent were to choose to put an extra load on the grid in order to fix their consumption, the humans would gain no extra utility from that load, and therefore nor should the home agent.

Each agent, i's utility is therefore given by:

$$U_i\left(\boldsymbol{x}, \boldsymbol{x}_a, \boldsymbol{\omega}, \ddot{\boldsymbol{\omega}}\right) = P_i\left(\boldsymbol{x}, \boldsymbol{x}_a, \boldsymbol{\omega}\right) + P_\omega\left(\omega_i, \ddot{\omega}_i\right) - f_r \cdot \omega_i - C_\theta\left(\theta_i, \alpha_{\theta,i}\right) - C_\gamma\left(\gamma_i, \alpha_{\gamma,i}\right) \tag{3.13}$$

where $\ddot{\boldsymbol{\omega}}$ is the vector of all agents' required loads.

Given these utility functions and the model presented in earlier sections, it is now possible to design mechanisms that elicit the desired behaviour from home agents. Analysis of the mechanisms presented in this thesis is both analytical and empirical. However, for empirical analysis, certain computational considerations must be taken into account, as is discussed next.

3.6 Computational Evaluation of the Model

While the above discussion presents the mathematical framework for the problem faced in this thesis which is sufficient for analytical analysis, empirical analysis must be performed through simulation on a computer. However, many aspects of the problem are computationally hard. For example, finding equilibrium strategies in games is known to be NP hard for normal form games (Conitzer and Sandholm, 2008), whereas in general, for games of over 4 players, finding optimal Nash equilibria is in the PPAD[2] class of problems (Daskalakis et al., 2006). Consequently, for computational simulation to be feasible over the large strategy space described in Equation 3.10, certain approximations and computational methods must be used. Thus, Section 3.6.1 describes a method of finding a (potentially non-optimal) Nash equilibrium known as iterated best response, and Section 3.6.2 discusses an approximation that allows for efficient computation of the sum of n log-normal distributions.

3.6.1 Finding Equilibria

It is not always possible to analytically find equilibria in games. As such, it is necessary when simulating games in computer systems to provide agents with a method for computationally finding equilibria within the game such that they are able to effectively strategise against one another. The problem of finding optimal equilibria within games is computationally hard (Daskalakis et al., 2006; Conitzer and Sandholm, 2008). Therefore, this section discusses a computational method for finding equilibria (which is not guaranteed to be optimal) – iterated best response. While, when iterated best response converges, it does so to a Bayes-Nash equilibrium, it is not guaranteed to converge to an equilibrium at all (Naroditskiy and Greenwald, 2007; Fudenberg, 1991).

The general algorithm for iterated best response is shown in Algorithm 1. In each round, k, each agent, i, knows what every other agent, $-i$, did in the previous round, which is represented by $\tilde{\boldsymbol{\sigma}}_i^{k-1}$. Given this, each agent chooses a strategy that maximises its expected utility given $\tilde{\boldsymbol{\sigma}}_i^{k-1}$ to play in this round. This is done using the optimisation algorithm described in Appendix D . At the end of every round, each agent's strategy is revealed to all agents, and the process is repeated until the strategies chosen by the agents converge to an equilibrium. This is equivalent to Cournot adjustment (Fudenberg and Levine, 1998). However it can result in cycles in which the strategies of agents do not

[2]Problems within PPAD are believed to be in NP, and therefore hard to compute (Daskalakis et al., 2009).

Algorithm 1 Pseudo-code for the iterated best response algorithm

$\tilde{\boldsymbol{\sigma}}_i^0 \leftarrow \left\{ \tilde{\sigma}_j^0 \middle| j \in N \setminus \{i\} \right\} \forall i \in N$ {Initialise the agents' strategy belief vector for all agents.}
$\bar{\boldsymbol{\sigma}}^0 \leftarrow \left\{ \tilde{\sigma}_i^0 = 0 \middle| \forall i \in N \right\}$ {A vector of agents' chosen strategies in the superscripted round}
$k \leftarrow 0$
while $\bar{\boldsymbol{\sigma}}$ not converged **do**
  $k \leftarrow k + 1$
  for $i \in N$ **do**
    $\bar{\sigma}_i^k = \arg\max_{\sigma_i' \in \Sigma_i} \bar{U} \left(\sigma_i' \middle| \tilde{\boldsymbol{\sigma}}_{-i}^{k-1} \right)$ {Agents choose their best response to their opponents strategies in the previous round}
  end for
  $\tilde{\boldsymbol{\sigma}}_i^k \leftarrow \left\{ \bar{\sigma}_j^k \middle| \forall j \in N \setminus \{i\} \right\} \forall i \in N$ {Agents are now able to see what each agent did in this round}
end while
The converged value of $\bar{\boldsymbol{\sigma}}$ (if any) is a Bayes-Nash equilibrium.

converge to a single strategy but rather a repeating sequence of strategies. To alleviate this problem, a technique called *partial best response* can be used whereby at any one time only a subset of all agents update their strategies based on previous iterations. However it was found that partial best response was unnecessary for the simulations within this thesis as iterated best response quickly converged successfully to equilibria.

3.6.2 Efficient Calculation of the Aggregate Belief

The general solution to the sum of a number of random variables is the convolution of those variables' probability density functions. However, this is computationally hard to compute, and as such, for the purpose of computational simulation, an approximation must be used. As the homes' consumptions are assumed to be log-normally distributed, constraints can be placed on their distributions stating that individual houses' expected consumptions are sufficiently far from zero, which allows one to use the approximate solution to the sum of their reports through using Gaussian distributions. Consequently, under empirical analysis, the aggregator calculates the aggregate consumption to be:

$$x^+ \sim \mathcal{N}\left(\mu^+, \theta^+\right)$$

where

$$\mu^+ = \sum_{i=1}^{n} \mathbb{E}\left(x_i\right)$$

and

$$\theta^{+} = \left[\sum_{i=1}^{n} \frac{1}{\theta_i} \right]^{-1}$$

Thus, with discussion of the model and the above computational methods complete, the next section will summarise what has been said and the remaining chapters will discuss the main contributions of the work in this thesis.

3.7 Summary

This chapter has provided an example scenario and a formalisation of the problem discussed within this thesis. Electricity market costs incurred by the aggregator were discussed in Section 3.2, along with how the aggregator's and home agents' beliefs are formed and how the aggregator determines the amount of electricity to purchase for the home agents. Section 3.3 discussed how agents may attempt to game the system, followed by a formal definition of agents' types and strategies in Section 3.4. Furthermore, the utility functions of both the aggregator and the home agents were discussed in Section 3.5. Finally, computational aspects of evaluating the model and any proposed mechanisms is discussed in Section 3.6.

The definition of the reward functions to be used by the mechanism is left to the next chapter. Specifically, Chapters 4 and 5 develop and provide both theoretical and empirical analysis of the *uniform* and *sum of others' plus max* mechanisms under the model described above.

Chapter 4

Mechanism Design for Scenarios with Report Manipulation

The model introduced in the previous chapter requires a payment function, P_i, which describes the amount to pay each agent based upon all agents' reports, the aggregator's priors, and the actual amount of electricity consumed by each agent. There are a number of ways of designing such a payment function. However, designing payments that elicit specific behaviours from agents in a way that is budget balanced is non-trivial. This chapter presents three payment mechanisms and discusses their theoretical properties. First, Section 4.1 presents the uniform mechanism; a Nash incentive compatible mechanism in which all agents are rewarded equally. Next, Section 4.2 introduces the percentage contribution mechanism. This mechanism, while being budget balanced, is in fact *not* incentive compatible – that is, agents are able to misreport in order to gain extra reward. Then, Section 4.3 discusses the sum of others' plus max mechanism; a *dominant strategy incentive compatible* mechanism which makes use of scoring rules to reward agents for their information in a way that is weakly budget balanced.

While sections 4.1—4.3 discuss the theoretical properties of the mechanisms presented in this chapter, empirical analysis is still required in order to determine the home agents' equilibrium behaviour under each of these mechanisms. That is, although it is known that agents will report truthfully under the uniform and sum of others' plus max mechanisms, agents are still able to strategise over the *actual amount of effort* they invest in building their estimates to report to the aggregator. Consequently, Section 4.4 analyses the empirical results in order to discuss how agents behave in equilibrium. Finally, Section 4.5 summarises the findings of this chapter.

4.1 The Uniform Mechanism

The simplest mechanism presented in this thesis is named the uniform mechanism and it simply divides the savings made by the aggregator equally amongst the agents. In this case, the reward given to each agent is:

$$P_i^{\mathbf{U}}(\boldsymbol{x}, \boldsymbol{x}_a, \boldsymbol{\omega}, n) = \frac{1}{n} \cdot \lambda \cdot \Delta(\boldsymbol{x}, \boldsymbol{x}_a, \boldsymbol{\omega}) \tag{4.1}$$

The uniform mechanism is budget balanced, as stated in Theorem 4.1 without the trivial proof. Moreover, the mechanism is *Nash* incentive compatible. That is, if an agent believes that all other agents will report truthfully, it will not gain a greater utility by misreporting. Theorem 4.2 provides a proof of Nash incentive compatibility.

Theorem 4.1. *The uniform mechanism is strictly budget balanced.*

Theorem 4.2. *The uniform mechanism is Nash incentive compatible, i.e. in terms of determining whether to misreport, for each agent $i \in N$, reporting $\hat{x}_i = x_i$ is a Nash equilibrium.*

Proof. When misreporting, the agents incur no other costs. Therefore, the agents' utilities are affected only by the reward they receive from the aggregator (as all other costs are 'sunk' – the cost of gathering the information to report to the aggregator has already been incurred by the agent). Therefore, each agent's expected reward is:

$$\mathbb{E}\left(P_i^{\mathbf{U}}(\boldsymbol{x}_a, \boldsymbol{x})\right) = \frac{1}{n} \cdot \lambda \cdot \bar{\Delta}(\boldsymbol{x}_a, \boldsymbol{x})$$

in which

$$\boldsymbol{x}_a = \boldsymbol{x}_{-i,a} \cup \{x_{i,a}\}$$
$$\boldsymbol{x} = \boldsymbol{x}_{-i} \cup \{\hat{x}_i\}$$

The expected savings function, derived from Equation 3.5, is taken to be the expected cost of the electricity given the aggregator's prior belief, $\bar{\kappa}(\boldsymbol{x}_a)$, minus the expected cost of the electricity given the agents' reports, $\bar{\kappa}(\boldsymbol{x})$, or:

$$\bar{\Delta}(\boldsymbol{x}_a, \boldsymbol{x}) = \bar{\kappa}(\boldsymbol{x}_a) - \bar{\kappa}(\boldsymbol{x})$$

From the agents' perspective, $\bar{\kappa}(\boldsymbol{x}_a)$ is constant. Furthermore, as a result of Equations 3.1 and 3.2 (that the aggregator purchases an amount which minimises the expected cost of electricity given the set of beliefs), the expected cost of electricity, $\bar{\kappa}(\boldsymbol{x})$, is *convex* with respect to $\boldsymbol{x}$. Consequently, any change in $\boldsymbol{x}$ from its true value (for example, from agent i misreporting $\hat{x}_i$) will *increase* $\bar{\kappa}(\boldsymbol{x})$, thereby *decreasing* savings and, as a consequence, agent i's expected reward. Thus, assuming all agents report truthfully, each agent maximises its utility by reporting truthfully, and any agent who unilaterally misreports will decrease its own reward. Therefore, truth telling is a Nash equilibrium. □

There is no guarantee that only one Nash equilibrium is present within a game. Indeed, under the uniform mechanism, not only is it a Nash equilibrium to report truthfully if all other agents do the same, it is also a Nash equilibrium to *misreport* if the other agents are expected to misreport themselves. This is shown in Theorem 4.3.

Theorem 4.3. *Under the uniform mechanism,* misreporting *is a Nash equilibrium.*

Proof. This proof follows a similar structure to that of Theorem 4.2. However it relies on the fact that:

$$\bar{\kappa}(\boldsymbol{x}) \equiv \bar{\kappa}\left(\mu^+, \theta^+\right)$$

where

$$\mu^+ = \sum_{x_i \in \boldsymbol{x}} \mu_i$$
$$\theta^+ = \left[\sum_{x_i \in \boldsymbol{x}} \frac{1}{\theta_i}\right]^{-1}$$

Let μ^+ and θ^+ be the respective aggregate mean and precision of the agents' *true* beliefs. Therefore, an aggregate belief of $\langle \mu^+, \theta^+ \rangle$ will *maximise* each agent's reward, as the aggregator purchases the amount of electricity that in expectation *minimises* the costs of electricity given $\langle \mu^+, \theta^+ \rangle$ (Equation 3.2). Now, suppose agent i reports a mean $\hat{\mu}_i = \mu_i + \varepsilon_i$, resulting in an aggregate belief of $\left\langle \hat{\mu}^+, \hat{\theta}^+ \right\rangle$, where $\hat{\mu}^+ = \mu^+ + \varepsilon_i$. It is now in the other agents' interest to *misreport*, in order to counteract the effect of agent i's misreport, and to thereby ensure that $\left\langle \hat{\mu}^+, \hat{\theta}^+ \right\rangle = \langle \mu^+, \theta^+ \rangle$. Thus, agent j reports a belief with mean $\hat{\mu}_j = \mu_j - \varepsilon_i$. The aggregate belief is now equal to the true belief, and thus all agents maximise their reward, and no single agent has an incentive to deviate. Hence, misreporting is a Nash equilibrium. □

The fact that misreporting is a Nash equilibrium is of course a huge disadvantage to the uniform mechanism – if an agent were to believe that its neighbours were to misreport, it will maximise its own reward by also misreporting. It could be argued that this form of 'counter-reporting' is advantageous to the system, or at least, provided the final aggregate belief is the same, it makes no difference to the system. However, computationally this is an undesirable situation as, for one, the agents have a larger state space over which to search for their optimal strategy. Furthermore, the agents may converge on a misreporting equilibrium, requiring extra computation in order to track the reports and expectations of other homes. This highlights another potential problem in the uniform mechanism, i.e. agents may be able to profit by their monitoring of one another. In many situations this may be acceptable, but within the smart grid domain this translates to neighbours gathering information on one-another's activities, which is clearly unacceptable.

These properties arise due to the fact that, using this mechanism, all agents are rewarded equally irrespective of their *actual* contribution. An ideal mechanism would reward the agents more *fairly*, by making greater payments to those agents whose estimates made the most significant increase in the aggregator's savings. A better solution would be a mechanism that is *dominant strategy incentive compatible.* That is, a mechanism wherein an agent's utility is maximised when reporting truthfully *regardless* of its belief of the other agents' actions. In the following sections, we demonstrate how this can be achieved using scoring rules. Scoring rules are specifically designed to rate predictors, as was discussed in Section 2.6. With this in mind, a scoring rule-based mechanism is developed in which an agent's reward is proportional to its score. However, in order to maintain budget balance, the scores received by the agents must somehow be scaled. As is shown in the next section, this is not trivial, as even an intuitively correct technique results in scores that are *not* incentive compatible.

4.2 The Percentage Contribution Mechanism

As was discussed in Section 2.6, scoring rules award a score to agents according to how well their prediction matches a given outcome. This makes scoring rules a natural tool for use in a mechanism which distributes payments based upon predictive information supplied by agents. Furthermore, given the requirement for *fair* distribution of the rewards to agents, it seems natural to distribute rewards proportionally to the amount they actually contributed to the overall result. The percentage contribution mechanism uses as a metric the ratio of the agent's score to the sum of *all* agents' scores within the system, as can be seen in Equation 4.2:

$$P_i^{\mathbf{P}}\left(\boldsymbol{x}, \boldsymbol{x}_a, \boldsymbol{\omega}\right) = \frac{S\left(\omega_i; \hat{\mu}_i, \hat{\theta}_i\right)}{\sum_{x_j \in \boldsymbol{x}} S\left(\omega_j; \hat{\mu}_j, \hat{\theta}_j\right)} \lambda \cdot \Delta\left(\boldsymbol{x}, \boldsymbol{x}_a, \boldsymbol{\omega}\right) \tag{4.2}$$

Insofar as the requirement for budget balance is concerned, Theorem 4.4 shows that the percentage contribution mechanism is *strictly* budget balanced. However, the mechanism is *not* incentive compatible, as discussed in Theorem 4.5.

Theorem 4.4. *The percentage contribution mechanism is strictly budget balanced.*

Proof. The total amount of savings distributed is:

$$\lambda \cdot \Delta\left(\boldsymbol{x}, \boldsymbol{x}_a, \boldsymbol{\omega}\right) \cdot \sum_{i \in N} \frac{S\left(\omega_i; \hat{\mu}_i, \hat{\theta}_i\right)}{\sum_{x_j \in \boldsymbol{x}} S\left(\omega_j; \hat{\mu}_j, \hat{\theta}_j\right)}$$

where the budget is $\lambda \cdot \Delta\left(\boldsymbol{x}, \boldsymbol{x}_a, \boldsymbol{\omega}\right)$. Thus it is sufficient to prove that the sum of the fraction is equal to one:

$$\begin{aligned}\sum_{i \in N} \frac{S\left(\omega_i; \hat{\mu}_i, \hat{\theta}_i\right)}{\sum_{x_j \in \boldsymbol{x}} S\left(\omega_j; \hat{\mu}_j, \hat{\theta}_j\right)} &\equiv \frac{1}{\sum_{x_j \in \boldsymbol{x}} S\left(\omega_i; \hat{\mu}_i, \hat{\theta}_i\right)} \sum_{i \in N} S\left(\omega_i; \hat{\mu}_i, \hat{\theta}_i\right) \\ &\equiv 1\end{aligned}$$

□

Theorem 4.5. *The percentage contribution mechanism is not incentive compatible.*

Proof. This is shown by counter-example. Figure 4.1 shows an agent's expected score when using the percentage contribution method of scaling with the spherical scoring rule. Specifically, Figure 4.1 shows an agent's expected score when it reports a belief with a scaled mean or variance. For each line, the other parameter is set to the agent's true value. For example, for the $\hat{\sigma}_i^2 = k \cdot \sigma_i^2$ line, the line shows the expected score for reporting $\hat{x}_i = \left\langle \mu_i, k \cdot \sigma_i^2 \right\rangle$. Thus, it can be seen from Figure 4.1 that the agent is able to increase its expected score by misreporting its precision. Therefore, the percentage contribution mechanism is not incentive compatible. □

Incentive compatibility is lost in the percentage contribution mechanism due to the interplay between agent i's score in the numerator and denominator of the fraction destroying the convexity of the expected score function. To address this issue, the sum of others' plus max mechanism was developed, which provides dominant strategy incentive compatibility. This mechanism and its theoretical properties are discussed in the next section.

4.3 The Sum of Others' plus Max Mechanism

As was discussed in the previous section, naively scaling scoring rules can lead to loss of incentive compatibility. However, this section presents a new scoring rule named *sum of others' plus max* which scales the scores received by agents such that they are still incentive compatible, and weakly budget balanced. This mechanism takes into account not only the spherical score achieved by the agent (defined in Table 4.1), but also those achieved by the other agents in the system. In the *sum of others' plus max mechanism*,

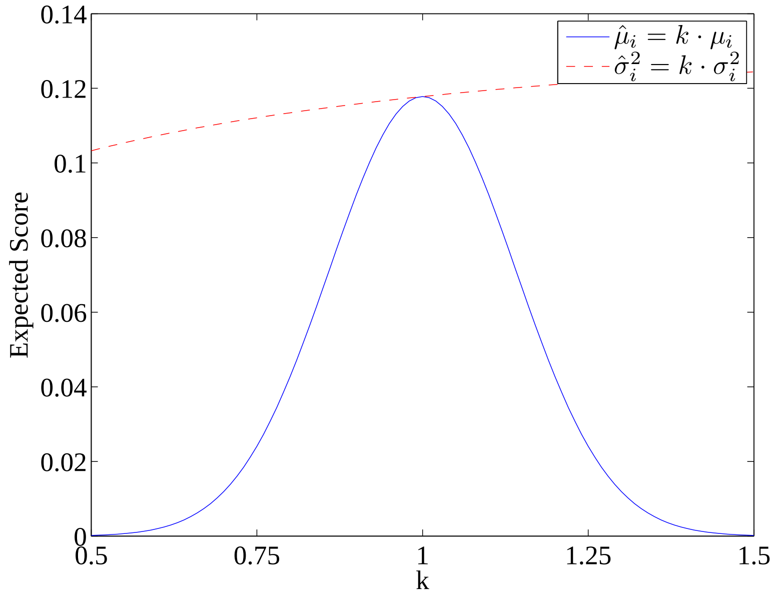

FIGURE 4.1: The expected score received by an agent who misreports its beliefs under the percentage contribution rule by scaling one parameter (either the mean or variance of their true belief) while reporting the other parameter truthfully.

payments are determined by multiplying the agents' scaled scores by the savings made by the aggregator when using the reports from the *other* agents in the system. This is necessary in order to preserve incentive compatibility. Furthermore, to ensure the agents' payments never outweigh the savings made by the aggregator, it is necessary to set a bound on the scores that are achievable by the home agents. A method of achieving this is discussed in the next section.

4.3.1 Bounding the Maximum Score

In order to use the sum of others' plus max mechanism, the scores achieved by the agents must be bounded. The logarithmic, quadratic and spherical rules are displayed in Table 4.1 in order to aid discussion.

The stationary points in Table 4.1 all represent points at which the score is *maximised*. It can be seen that the upper bound of these scores are all determined by the probability density function of the belief which they are scoring. Consequently, in order to place an upper bound on the scores assigned by these rules, constraints must be placed upon the acceptable beliefs. This section discusses the methods of achieving this for the commonly used Gaussian distribution, and the log-normal distribution, which is used by the agents in the simulations within this thesis.

	Logarithmic Rule	Quadratic Rule	Spherical Rule
$S(x,\omega)=$	$\log\left(x\left(\omega\right)\right)$	$2x\left(\omega\right)-\displaystyle\int x\left(\omega'\right)^2 \mathrm{d}\omega'$	$\dfrac{x\left(\omega\right)}{\sqrt{\int x\left(\omega'\right)^2 \mathrm{d}\omega'}}$
$\dfrac{\mathrm{d}}{\mathrm{d}\omega}S(x,\omega)=$	$\dfrac{\mathrm{d}x\left(\omega\right)}{\mathrm{d}\omega}\cdot\dfrac{1}{x\left(\omega\right)}$	$2\cdot\dfrac{\mathrm{d}x\left(\omega\right)}{\mathrm{d}\omega}$	$\dfrac{\mathrm{d}x\left(\omega\right)}{\mathrm{d}\omega}\cdot\dfrac{1}{\sqrt{\int x\left(\omega'\right)^2 \mathrm{d}\omega'}}$
SPs:	$\dfrac{\mathrm{d}x\left(\omega\right)}{\mathrm{d}\omega}=0,\ x\left(\omega\right)=\infty$	$\dfrac{\mathrm{d}x\left(\omega\right)}{\mathrm{d}\omega}=0$	$\dfrac{\mathrm{d}x\left(\omega\right)}{\mathrm{d}\omega}=0$

Table 4.1: A comparison of three major scoring rules within the literature showing their first derivatives and stationary points (SPs) with respect to the realised event, ω. The term, $x\left(\omega\right)$, represents the probability density at event ω.

4.3.1.1 Gaussian Distributions

The Gaussian distribution is fairly simple to bound. Given a mean, μ, and standard deviation, σ, the maximum probability density, $\rho_{max}\left(\mu,\sigma\right)$, can be calculated using:

$$\rho_{max}\left(\mu,\sigma\right)=\mathcal{N}\left(\omega^*;\mu,\sigma\right)$$

where

$$\omega^*=\arg\max_{\omega}\mathcal{N}\left(\omega;\mu,\sigma\right) \tag{4.3}$$

thus

$$\omega^*=\mu \tag{4.4}$$

$$\rho_{max}\left(\mu,\sigma\right)=\frac{1}{\sigma\sqrt{2\pi}} \tag{4.5}$$

Equation 4.5 is also plotted in Figure 4.2. Both Figure 4.2 and Equation 4.5 show that, in order to restrict the upper bound of the probability density of an event, the standard deviation must itself have an *lower bound* enforced. Recall that precision of a distribution is one over its *variance*. Thus, placing a lower bound on the standard deviation is equivalent to placing an *upper* bound on precision.

Thus, in order to use the sum of others' plus max mechanism with *Gaussian* distributions, the precision of the agents' reports must be bounded. This seems counter intuitive as the aggregator benefits from more precise estimates. However, in reality, the bound on precision can be set arbitrarily high. For example, the aggregator may set the precision

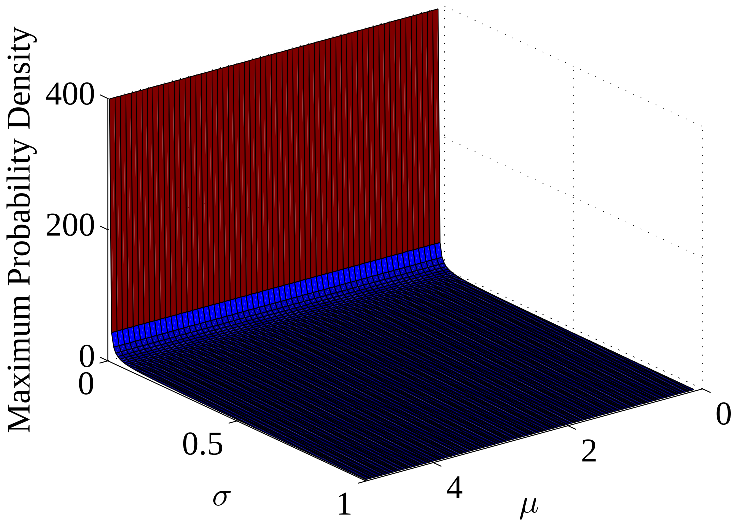

FIGURE 4.2: The maximum probability density given by the Gaussian probability density function for means, μ, between zero and four, and standard deviations, σ, between zero and one.

to a value so high that it is unlikely any of the home agents' reports will be restricted by it.

While this discussion provides an example of restricting the probability density of the home agents' beliefs, assuming those beliefs are Gaussian, in the simulations presented within this work, log-normal distributions are used. Constraining the probability density follows a similar procedure as shown in this section and it is shown that once again, constraining the variance is sufficient to place an upper bound.

4.3.1.2 Log-Normal Distributions

Unlike Gaussian distributions, the peak probability density of a log-normal distribution does not necessarily occur at the distribution's expected value. In fact, the peak probability density occurs at:

$$\omega^* (\mu, \sigma) = \arg\max_{\omega} \log \mathcal{N} (\omega; \mu, \sigma)$$

or

$$\omega^* (\mu, \sigma) = \mathrm{e}^{\mu - \sigma^2} \tag{4.6}$$

Furthermore, by substituting Equation 4.6 into the probability density function for the log-normal distribution, the upper bound on the probability density function given any

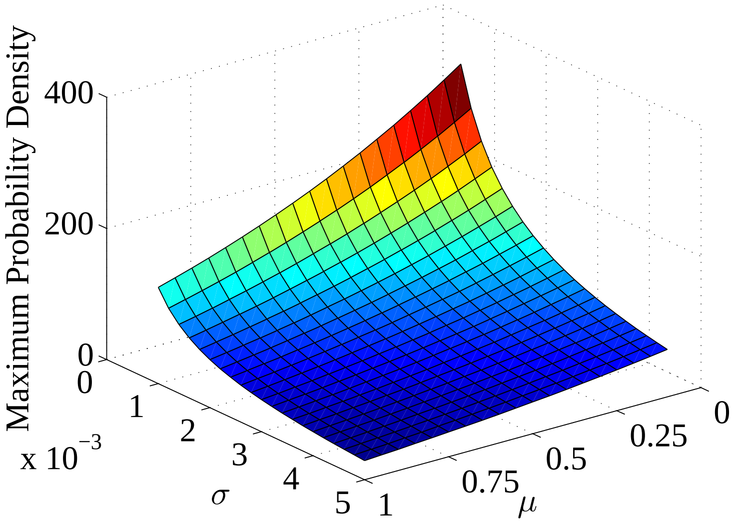

FIGURE 4.3: The maximum probability density given by the log-normal probability density function for means, μ, between 0 and 1, and standard deviations, σ, between 0 and 5×10^{-3}.

mean, μ and standard deviation, σ can be found by[1]:

$$\begin{aligned} \rho_{max}(\mu, \sigma) &= \log \mathcal{N}(\omega^*(\mu, \sigma); \mu, \sigma) \\ &= \frac{e^{-\mu+\frac{\sigma^2}{2}}}{\sqrt{2\pi}\sigma} \end{aligned} \tag{4.7}$$

Equation 4.7 is plotted in Figure 4.3, which shows that the probability density of a log-normal distribution only becomes large as *both* the mean and standard deviation approach zero. In fact, as $\mu \to 0$, the maximum probability density becomes dependent on the standard deviation. Furthermore, as is the case with the Gaussian distribution, taking the limit of Equation 4.7 as $\sigma \to 0$ results in infinity. That is:

$$\lim_{\mu \to 0} \frac{e^{-\mu+\frac{\sigma^2}{2}}}{\sqrt{2\pi}\sigma} = \frac{e^{\frac{\sigma^2}{2}}}{\sqrt{2\pi}\sigma}$$

and:

$$\lim_{\sigma \to 0} \frac{e^{-\mu+\frac{\sigma^2}{2}}}{\sqrt{2\pi}\sigma} = \infty$$

Consequently, in order to constrain the maximum probability density of any particular event in a log-normal distribution, it is necessary only to bound the *standard deviation*,

[1]Although the log-normal distribution isn't strictly parameterised by its mean and standard deviation, it is possible to use them through mapping them to the *actual* parameters using equations 4.10 and 4.11.

or *precision*, as was the case with the Gaussian distribution.

4.3.2 Properties of the Sum of Others' plus Max Mechanism

Given the discussion in the previous section, an upper bound is placed upon the precision of the reports on which the home agents will be scored. Let this maximum precision be denoted by θ_{max}. If any agent reports a precision $\hat{\theta}_i > \theta_{max}$, their spherical score will be calculated as though $\hat{\theta}_i = \theta_{max}$. Formally, with this substitution performed, the payment agent i obtains, given all agents' estimates is given by:

$$P_i^{\mathbf{S}}\left(\boldsymbol{x}, \boldsymbol{x}_a, \boldsymbol{\omega}\right) = \frac{S\left(\omega_i; \hat{\mu}_i, \hat{\theta}_i\right) \cdot \lambda \cdot \Delta\left(\boldsymbol{x}_{-i} \cup \{x_{a,i}\}, \boldsymbol{x}_a, \boldsymbol{\omega}\right)}{S\left(\omega_j; \omega_j, \theta_{max}\right) + \sum_{x_j \in \boldsymbol{x}_{-i}} S\left(\omega_j; \hat{\mu}_j, \hat{\theta}_j\right)} \tag{4.8}$$

Here, assuming a Gaussian distribution, the $S\left(\omega_j; \omega_j, \theta_{max}\right)$ term represents the maximum score that can be achieved by an agent – the score achieved when reporting the maximum possible precision, θ_{max}, and reporting an estimate with mean ω_j when ω_j actually does occur. It can be seen that, by using only the savings made by the other agents, the only term that is dependent on the report from the agent being rewarded is the scoring rule. Moreover, the spherical scoring rule was specifically chosen due to its strict propriety, and therefore these payments are incentive compatible, as is shown in Theorem 4.6.

Theorem 4.6. *The sum of others' plus max mechanism is strictly dominant strategy incentive compatible, while $\theta_i < \theta_{max}$, for any strictly proper scoring rule, S. i.e. truth telling is a strictly dominant strategy.*

Proof. The maximum score is a constant value set by the mechanism designer. Furthermore, the agent's report is excluded from the calculation of the savings made. This results in the agent being unable to affect the savings made by the other agents. Consequently, the savings made by the other agents are in effect a constant. Given this, sum of others' plus max is simply an affine transformation of the spherical scoring rule, which maintains strict propriety and therefore incentive compatibility. The fact that the score is *strictly* proper means that the expected score has a unique maximum where an agent reports truthfully. Therefore, the expected reward an agent receives is also a unique maximum when the agent reports truthfully, and thus the mechanism is strictly dominant incentive compatible. □

In addition to truth telling being a strictly dominant strategy, the rewards to agents using this rule are fairer than those of the uniform mechanism in that the agents are rewarded based upon their *own* reports. Consequently, an agent who puts in a small amount of effort to generate a low precision estimate will expect to receive a small reward whereas if the agent were to put in a large amount of effort, the agent would score more highly and therefore receive a larger reward.

In order to maintain weak budget balance, it is essential to divide the agent's spherical score by the sum of the other agents' prescaled scores *plus the maximum score*. This ensures that the fraction is *always* less than one, and always sums to *at most one*. This fact is used in Theorem 4.7, which provides a proof of the fact that sum of others' plus max is ex ante weakly budget balanced.

Theorem 4.7. *The sum of others' plus max mechanism is ex ante weakly budget balanced.*

Proof. Let each agent, i, obtain the score S_i, and let $\bar{\Delta}(\theta)$ be the expected savings made when the agents' reports produce an aggregate precision of θ. Under sum of others' plus max, when each agent, i, generates information of precision θ_i, the total expected payment is:

$$\sum_{\forall i \in N} \frac{S_i}{S_{max} + \sum_{\forall j \in N \setminus \{i\}} S_j} \cdot \lambda \cdot \bar{\Delta}\left(\left(\sum_{\forall j \in N \setminus \{i\}} \frac{1}{\theta_j} + \frac{1}{\theta_{a,i}}\right)^{-1}\right)$$

and the aggregator's total expected savings that is allocated for rewarding the agents is

$$\lambda \cdot \bar{\Delta}\left(\left(\sum_{\forall i \in N} \frac{1}{\theta_i}\right)^{-1}\right)$$

The sum of the fraction of scores is ≤ 1, and $\bar{\Delta}(\theta)$ is strictly increasing with θ. Therefore, it is sufficient to prove:

$$\left(\sum_{\forall j \in N \setminus \{i\}} \frac{1}{\theta_j} + \frac{1}{\theta_{a,i}}\right)^{-1} \leq \left(\sum_{\forall i \in N} \frac{1}{\theta_i}\right)^{-1} \quad \forall i \in N$$

Start with the fact from Section 3.2 that the aggregator's precision is less than or equal to the agent's:

$$\theta_{a,i} \leq \theta_i$$

Adding $\theta_i \theta_{a,i} \gamma > 0$ to both sides gives:

$$\theta_{a,i} + \theta_i \theta_{a,i} \gamma \leq \theta_i + \theta_i \theta_{a,i} \gamma$$

Which factorises to give:

$$\theta_{a,i}\left(\theta_i\gamma+1\right)\leq\theta_i\left(\theta_{a,i}\gamma+1\right)$$

The bracketed expressions are strictly positive. Therefore, it can be simplified to give:

$$\left(\gamma+\frac{1}{\theta_{a,i}}\right)^{-1}\leq\left(\gamma+\frac{1}{\theta_i}\right)^{-1}$$

Substituting γ for $\sum_{\forall j\in N\setminus\{i\}}\frac{1}{\theta_j}$, we are left with

$$\left(\sum_{\forall j\in N\setminus\{i\}}\frac{1}{\theta_j}+\frac{1}{\theta_{a,i}}\right)^{-1}\leq\left(\sum_{\forall i\in N}\frac{1}{\theta_i}\right)^{-1}$$

That is, the precision of the aggregate report made by the aggregator when using its own information in place of agent i's is less than the precision of using all agents reports, when each agent reports a precision greater than or equal to the aggregator's precision. Combined with the fact that the expected savings are strictly increasing with precision, and the sum of fractions of the budget allocated to each agent is less than or equal to one, this shows the sum of others' plus max mechanism is *ex ante weakly* budget balanced. □

There are numerous advantages to using sum of others' plus max over the simple uniform mechanism presented earlier. Firstly, truth telling strictly dominates all other strategies. As a result, reporting truthfully will *always* maximise the agent's expected reward, regardless of the other agents' actions. This is not the case in the uniform mechanism wherein truth telling is only a Nash equilibrium. For example, if an agent learns that its neighbour is to misreport its estimate, it too could misreport in order to offset the other agent. However, a disadvantage of sum of others' plus max compared to the uniform mechanism is that it is only *ex ante* weakly budget balanced – a weaker concept than the ex post strict budget balance exhibited by the uniform mechanism. This is further explained with the aid of empirical evidence in Section 4.4.6. The home agents, might also make small losses when the other agents' predictions are poor. However, in expectation, home agents' utilities will always be positive as they are able to strategise over the precision the generate in order to maximise their utility.

4.4 Empirical Evaluation of the Mechanisms

The previous sections discussed the theoretical properties of the mechanisms presented within this thesis. In doing so, the behaviours adopted by agents have only been discussed in terms of truth telling. However, it is not clear what the precision of the reports

Symbol	Description	Value
—	Trials	1000
n	Number of agents	2–16
θ_{max}	Maximum scored precision	1500
λ	Fraction of savings to distribute	0.5
f	Forward market price	100
δ^b	System buy price delta	50
δ^s	System sell price delta	70
f_r	Retail price per unit of electricity that the agents must pay to the aggregator	150
$\alpha_{\theta,i}$	Agent i's precision cost coefficient	$i/100$
$\alpha_{\gamma,i}$	Agent i's load variance manipulation cost coefficient (Not applicable to this scenario)	—

Table 4.2: The values of parameters used by the simulation for Chapter 4.

generated by those agents will be, and their effect on the utility to the aggregator. To address this issue, this section presents an empirical analysis of the mechanisms. In more detail, the equilibrium strategies under each mechanism is found using the process of iterated best response described in Section 3.6.1. Using simulation, these strategies are evaluated to determine, the utilities to the agents and aggregator.

The remainder of this section is structured as follows: Section 4.4.1 discusses the configuration of the simulation; Section 4.4.3 discusses on the agents' behaviour under 'no mechanism', i.e. when the aggregator uses its priors and does not request information from the home agents nor reward them; Section 4.4.4 discusses the precision of reports offered by the home agents under the mechanisms presented above; Section 4.4.5 discusses the utility obtained by the agents and aggregator in using the uniform and sum of others' plus max mechanisms over using no mechanism; Section 4.4.6 presents the risk to the aggregator under these various mechanisms and, finally Section 4.4.7 presents a related discussion on the improvement of social welfare through the use of the mechanisms.

4.4.1 Simulation Setup

This section provides a description of the setup of the simulations run in order to provide the empirical results that are analysed within this chapter. For quick reference, Table 4.2 provides an overview of the values of the various model parameters as used in the simulations.

4.4.1.1 UK Residential Consumption Data

The consumptions of homes within the simulations are not sampled from arbitrary distributions. Rather, they are based on real consumption data from residential properties

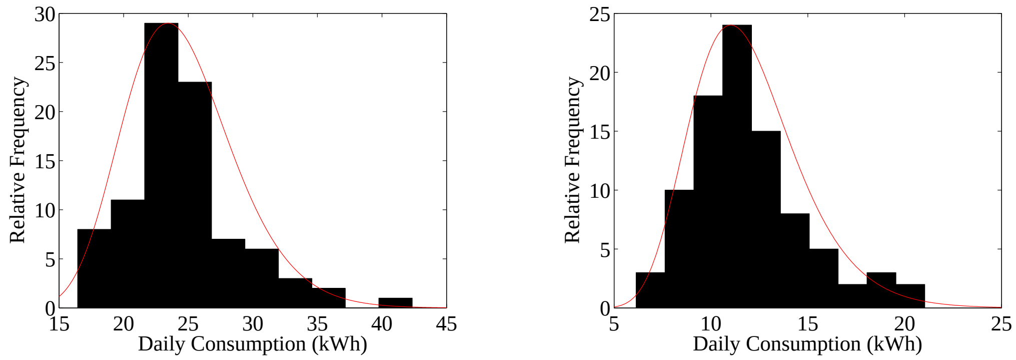

FIGURE 4.4: Examples of the distribution of daily consumptions from two of the houses used within the simulations in this thesis. Histograms of the raw data are shown in black, with a red line representing the scaled probability density function of a fitted log-normal distribution.

within the UK, which was obtained from a Hampshire-based smart metering company. The data contains half-hourly consumptions for a set of homes within the UK. However, almost all homes contain missing data. To resolve this, houses with large numbers of missing data points are rejected. From the remaining houses, for each house containing missing data points, a log-normal distribution is fitted to the existing data for that house, and the missing points were sampled from the resulting distribution.

The half-hourly data points correspond to the consumption of each property within each half-hour trading slot within the electricity market. However, in order to ensure the house consumptions are far from zero (as per the discussion in Section 3.6.2), the half-hourly data points were summed to calculate the total consumption for each house over 24 hour periods. This results in a total of 80 data points per house. Therefore, houses within the simulation are, in effect, asked to predict their *daily* consumption. It should be noted that this does not affect the generality of the results in this thesis. The mechanism simply requires the home agents to make predictions for a specified time period and is indifferent as to the period's actual length. Examples of the resulting distributions of houses consumptions are shown in Figure 4.4.

As can be seen from Figure 4.4, log-normal distributions appear to fit the data fairly well. Furthermore, log-normal distributions implicitly represent the fact that the homes' consumptions must be strictly positive whilst still allowing beliefs of low precision with expected values close to zero. If the more mathematically convenient Gaussian distribution were used, care would have had to be given to ensure that the distributions' expected values were sufficiently far from zero such that the tail of the distribution that resided in negative space would contribute a negligible amount to any calculations. Of course, Figure 4.4 shows just 2 of the 16 total homes used within this thesis. Plots of all 16 homes' consumptions can be seen in Appendix C. Given this data, the next section describes how agents' reports are generated.

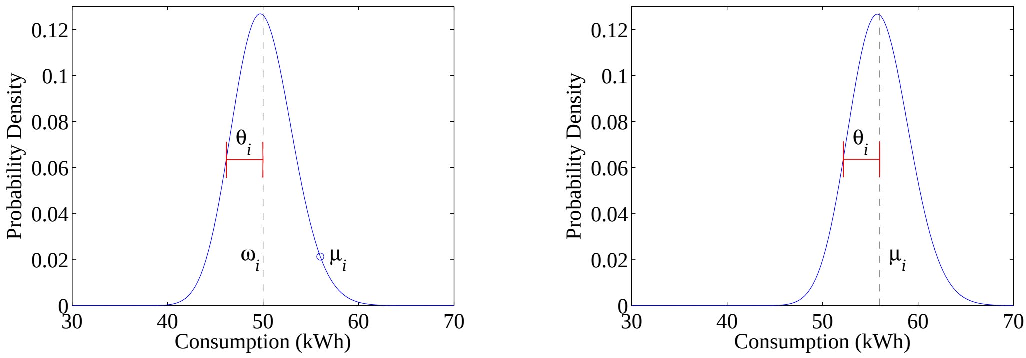

(a) Distribution from which μ_i is sampled.

(b) Agent i's belief, $x_i = \langle \mu_i, \theta_i \rangle$.

FIGURE 4.5: A diagram showing how the agent's belief's mean is determined based upon that agent's future consumption, ω_i, and the agent's chosen precision, θ_i, (a), and agent i's resulting belief, (b)

4.4.1.2 Simulating Agents' Reports

In simulation, the agents must generate reports that represent some imprecise prediction of the agents' future consumption based on collected information. This process is simulated by generating noisy predictions based upon the agent's future consumption, which is done as follows:

Each agent i's consumption, ω_i, is sampled from that home's particular dataset. Given an agent's chosen precision, θ_i, the *expected value* of the agent's report, μ_i is calculated as shown graphically in Figure 4.5, or more formally, as follows:

$$\mu_i \sim \log \mathcal{N}\left(\mu_{L,i}, \sigma^2_{L,i}\right) \tag{4.9}$$

where

$$\mu_{L,i} = \log\left(\omega_i\right) - \frac{1}{2}\log\left(1 + \frac{1}{\theta_i \cdot \omega_i^2}\right) \tag{4.10}$$

$$\sigma^2_{L,i} = \log\left(1 + \frac{1}{\theta_i \cdot \omega_i^2}\right) \tag{4.11}$$

Agent i's belief then becomes $x_i = \langle \mu_i, \theta_i \rangle$. In this way, agents' beliefs represent a noisy view of their future consumption. Moreover, as an agent invests greater effort into producing reports with greater information, the expected value of their belief approaches their actual consumption.

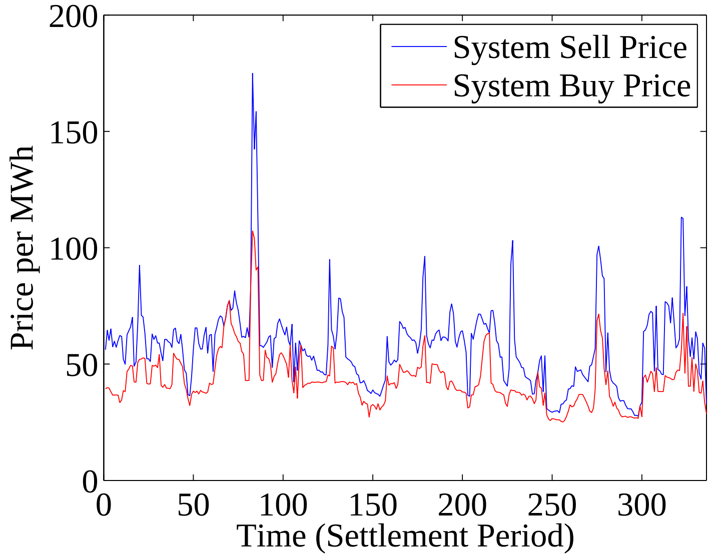

FIGURE 4.6: System buy and sell prices for the 22nd January 2013 to the 28th January 2013 inclusive.

4.4.1.3 Market Parameters and Agents' Costs

The values for market prices are set such that it is more expensive for an agent to over-consume than to under-consume. This represents the fact that extra generation would have to be supplied to fill the unexpected demand, and encourages more conservative use of electricity. Due to the unit-agnostic approach taken to the simulations here, a central price of 100 was chosen as the forward market price, f. Heavy penalties are imposed on over-consuming, with a system sell price delta, $\delta^s = 70$. Lower penalties are imposed on those under-consuming, setting the system buy price delta, $\delta^b = 50$. Given these penalties, the aggregator sells the electricity to the consumers at a price of $f_r = 150$.

It is difficult to measure how realistic these parameters are. For one, the analysis performed in this thesis is unit agnostic in that there is no specific unit of currency, and the forward market is simply set to a central price of 100. Moreover, the mechanism here is concerned only with a single timeslot, and consequently uses the fixed values discussed above. However, in real life, the balancing prices are heavily influenced by the overall demand on the grid and vary constantly over time. As an example, Figure 4.6 shows the real balancing market prices for a week between the 22nd January 2013 to the 28th January 2013 inclusive.

Furthermore, it is difficult to ascertain how realistic the balancing market prices are with respect to the forward market prices. While in real life the balancing market prices are dependent upon the spot-market prices, electricity companies are able to procure electricity in the form of long-term contracts which will affect the exact cost per kilowatt hour that the electricity companies incur. With this in mind, the empirical results within this thesis should not be read in order to determine absolute values of utility. Instead, they should be used only to compare the various mechanisms that are tested. That is, it should be said that one mechanism performs better than the other rather than that under one mechanism the aggregator makes $\$x$ whereas under the other mechanism the aggregator makes $\$y$.

In terms of the home agents, the cost of generating a report of a given precision can arise from many factors. For example, in order to make more precise information, the home owners have to invest more time in adding data-sources to their smart meters, or extra computational power might be used in generating these reports. However, these costs are typically relatively small when compared to the cost of electricity per unit (not least because the total cost is divided over all units of electricity purchased). Moreover, it is likely that different homes have different costs. For example, the occupants of one house might value their time more highly than their neighbours. Alternatively, differing hardware can result in different costs. For example, the sensors in one home might be more modern and less power-hungry than the sensors in another home. Consequently, even assuming the same amount of computational effort[2] in generating the information for the aggregator, the house with the older sensors will consume more electricity when generating their information. Thus, with this in mind each agent i's precision coefficient (the cost per precision, θ_i, that agent i incurs) is taken to be an arbitrarily small value, $\alpha_{\theta,i} = i/100$.

Now that the basic setup of the simulations has been described, all that remains is to describe a benchmark against which the uniform and sum of others' plus max mechanisms will be compared. This is the scenario in which agents do not report any information and the aggregator must rely on its own priors. This is referred to as the 'no mechanism' scenario, which is discussed in more detail next.

4.4.1.4 The Allocation of the Budget

The aggregator does not necessarily distribute all of the savings it makes by using the agents' information. The λ parameter determines the portion of the savings the aggregator actually allocates to the agents. The actual value of this parameter affects how the agents behave as well as the utility to the home agents. Figure 4.7(a) shows the aggregator's expected utility (minus the money made by selling the electricity to the agents) for differing λ in a system of 20 agents, when the parameters are set as described in Table 4.2. It can be seen that for $\lambda = 0$ the aggregator's expected utility is zero (within errors). This comes about due to the fact that even when lambda is zero, and consequently the aggregator is not distributing any money to the home agents, still sometimes the aggregator might make a loss due to its own prediction error. It can be seen that setting $\lambda = 0.2$ gives a much higher utility for the aggregator, as agents are

[2]Of course, the more modern sensor might have optimised hardware which allows it to more computationally efficiently generate the information.

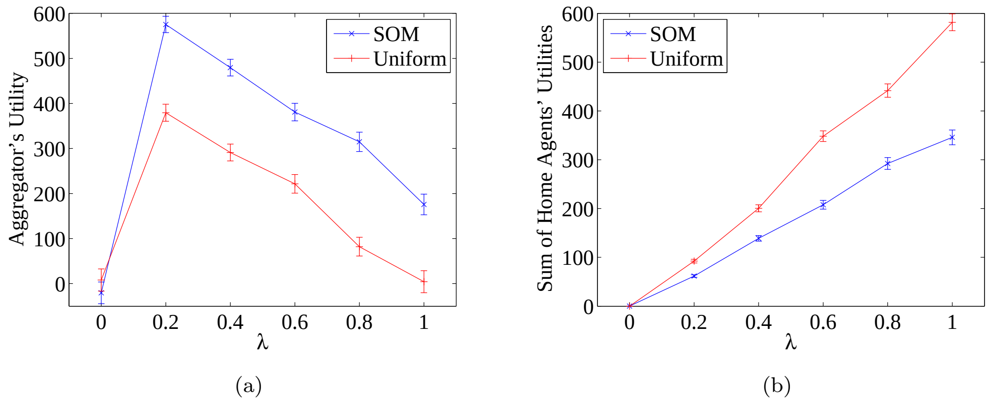

FIGURE 4.7: The effect of varying the fraction of savings to reward agents, λ on: (a), the aggregator's expected utility; (b) on the home agents' expected utility for 20 agents

encouraged to make predictions for the aggregator. Furthermore, it can be seen that as $\lambda \to 1$, the utility of the aggregator gradually reduces as the aggregator is giving away more of its savings.

Interestingly, it can be seen that at $\lambda = 1$, the aggregator still makes a positive utility from the sum of others' plus max mechanism whereas under the uniform mechanism, the aggregator gains zero utility. This is due to the *weakly* budget balanced nature of sum of others' plus max versus the *strict* budget balanced nature of the uniform mechanism. At $\lambda = 1$, under the uniform mechanism, the aggregator gives away all of the savings made by the agents. Consequently, the aggregator gains no extra utility. However, under sum of others' plus max, when $\lambda = 1$ the aggregator gives away *at most* all of its savings. The remaining undistributed savings are kept by the aggregator, giving it positive utility.

Conversely, it can be seen that the home agent's utilities, the sum of which is shown in Figure 4.7(b), increases as $\lambda \to 1$. It can be seen that at $\lambda = 0$ the home agents gain no utility. This is of course because there is no budget for allocating a reward to the home agents. More importantly, it can be seen that when $\lambda = 0$, the agents do not make a loss. This shows that when λ is low, agents choose not to invest any effort in generating reports, as doing so would not be rational.

Against this, for the simulations in the remainder of this chapter, which test the equilibrium strategies of the home agents against the number of home agents, the value $\lambda = 0.5$ is used. Due to the fact that the aggregator loses utility as $\lambda \to 1$ and the home agents lose utility as $\lambda \to 0$, this seems like a reasonable compromise such that exactly half of the savings are allocated to the budget for rewarding the agents and half is guaranteed to be retained by the aggregator.

4.4.2 Simulation Methodology

Each simulation performs 1000 repeats of an entire game. That is, each agent's consumption is randomly selected from that agent's data (discussed above in Section 4.4.1.1), their initial strategies are randomised and then iterated best response is performed in order to find each agent's equilibrium strategy, as was discussed in Section 3.6.1. Randomising the starting strategies in this way means that, if multiple equilibria exist, they are more likely to be found. However, as can be seen later on when discussing the agents' precisions, even when the agents' starting strategies are randomised, iterated best response converged to the same equilibrium strategies.

Once the equilibrium strategies are found, the process of generating the agents' reports, reporting them to the aggregator and then the aggregator purchasing electricity and rewarding the agents is performed. The strategies and resulting utilities of the home agents and the aggregator are stored for each of the 1000 repeats. This raw data is then used to generate the plots shown in the upcoming sections. The means shown in these plots are taken over the 1000 repeats, and error bars show standard error calculated as $\sqrt{1000/\sigma^2}$, where σ^2 is the variance of the raw data being plotted.

4.4.3 The 'No Mechanism' Scenario

Many of the results in the following sections are compared to a situation in which no mechanism is implemented by the aggregator. This corresponds to the current situation within the electricity markets. That is, the aggregator must purchase electricity for the consumers using *only* its prior information. Therefore, under such a situation, home agents do not gather any information and nor do they incur any costs. Nor do agents receive any payment from the aggregator. Consequently, this scenario represents a natural benchmark on which to compare the mechanisms developed in this thesis. In particular, the aggregator's expected utility should be greater under the developed mechanisms than in the no mechanism case. If this is not the case, it is not rational for the aggregator to implement the given mechanism as doing so will damage its own utility.

Moreover, home agents should receive a higher utility under the mechanisms developed within this thesis than they would under then 'no mechanism' scenario. In the contrasting case, in which the mechanism is not individually rational, the agents will simply choose not to perform any data collection and not to report any information to the aggregator. Consequently, a non-individually rational mechanism will result in agents behaving as though there is no mechanism, thereby not incurring any costs. However, by the agents exhibiting this behaviour, the aggregator gains no extra utility by implementing the non-individually rational mechanism and thus may as well implement the 'no mechanism' scenario.

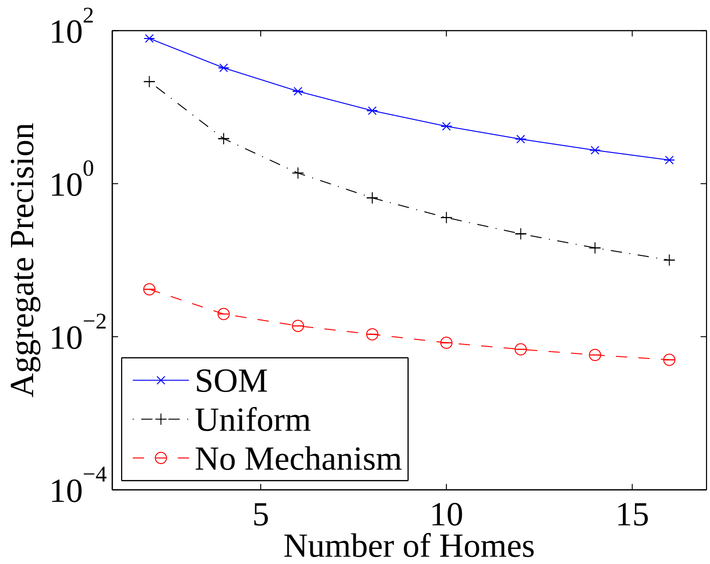

Figure 4.8: A figure showing the precision of the aggregate belief used by the aggregator in buying electricity for the homes.

Ideally, a mechanism will give agents positive utility such that they are encouraged to make precise reports. The exact precision of the reports generated by the agents under the uniform and sum of others' plus max mechanisms are discussed in the next section.

4.4.4 Precision of Reports

An important property of the mechanisms implemented by the aggregator is that they encourage precise prediction of consumptions from the home agents. Increasing the precision of predictions from the home agents will allow for more precise prediction of the aggregate consumption, more efficient scheduling of generation, and a reduced cost of electricity to the aggregator. This section discusses the precision of the aggregate report used by the aggregator in order to purchase electricity under the 'no mechanism' scenario, the uniform mechanism and the sum of others' plus max mechanism.

The 'no mechanism' plot from Figure 4.8 represents the precision of the aggregate belief when using *only* the aggregator's priors. Thus, it can be seen that both the uniform and the sum of others' plus max mechanisms encourage agents to provide reports with greater precision than the aggregator's priors. It can be seen that the prior information held by the aggregator is imprecise in comparison to the information received from the agents in the uniform and sum of others' plus max mechanism. However, note that this value is *not* zero. It can be seen that the sum of others' plus max mechanism offers the best increase in the aggregate precision. Furthermore, it can be seen that as the number of agents increases, the agents' aggregate precision decays towards that of the aggregator's prior information. This is due to a number of factors. Primarily, as the number of agents increases, the savings made by the aggregator *decreases*. This results from the fact that prediction errors are more likely to be cancelled by an error in the opposite direction for another home for larger numbers of homes. As such, the cost of electricity when using the aggregator's priors decreases as the number of homes increases. Consequently, the savings decreases.

These precisions come about *not* due to the absolute value of the rewards the agents receive but rather the *gradient* of the rewards with respect to θ_i. That is, each agent i chooses its precision, θ_i by solving the first derivative of its utility function (from Equation 3.13) for:

$$\frac{\mathrm{d}}{\mathrm{d}\theta_i} U_i\left(\boldsymbol{x}, \boldsymbol{x}_a, \boldsymbol{\omega}, \ddot{\boldsymbol{\omega}}\right) = 0 \tag{4.12}$$

or

$$\frac{\mathrm{d}}{\mathrm{d}\theta_i} P_i\left(\boldsymbol{x}, \boldsymbol{x}_a, \boldsymbol{\omega}\right) = \alpha_{\gamma,i}. \tag{4.13}$$

The only time the *absolute* reward is of importance is in determining whether a rational agent would make a report or not. It can be seen in the next section that while the agents in the sum of others' plus max mechanism invest a greater utility in generating their reports from the sum of others' plus max mechanism, they gain a *smaller* amount of utility than they do in the uniform mechanism; further evidence that it is the *gradient* of the agents' utility that is important rather than the *absolute* value of utility.

4.4.5 Utility to the Home Agents

While the home agents have no choice in the mechanism being used, it is still beneficial to analyse the utility received by the agents under equilibrium. In particular, the mechanisms should be *individually rational*. That is, the home agents should not expect to receive negative utility.

Figure 4.9 shows how the average utility per agent changes as the number of agents increases. The fact that both the uniform and sum of others' plus max lines are positive shows that the mechanism is individually rational for the home agents (that is, agents expect to make a positive utility). However, it shows that in terms of the home agents' utility, the uniform mechanism outperforms the sum of others' plus max mechanism. This is due to the fact that the home agents are investing a far greater amount of effort in the sum of others' plus max mechanism (as can be seen from Figure 4.8), combined with the fact that the sum of others' plus max is *weakly* budget balanced and thus not all of the allotted budget is rewarded to the home agents.

While giving a large amount of utility to the home agents is by no means a bad thing, the actual amount the aggregator pays to the home agents is not necessarily of large importance provided that the agents receive positive utility. That is, the fact that the

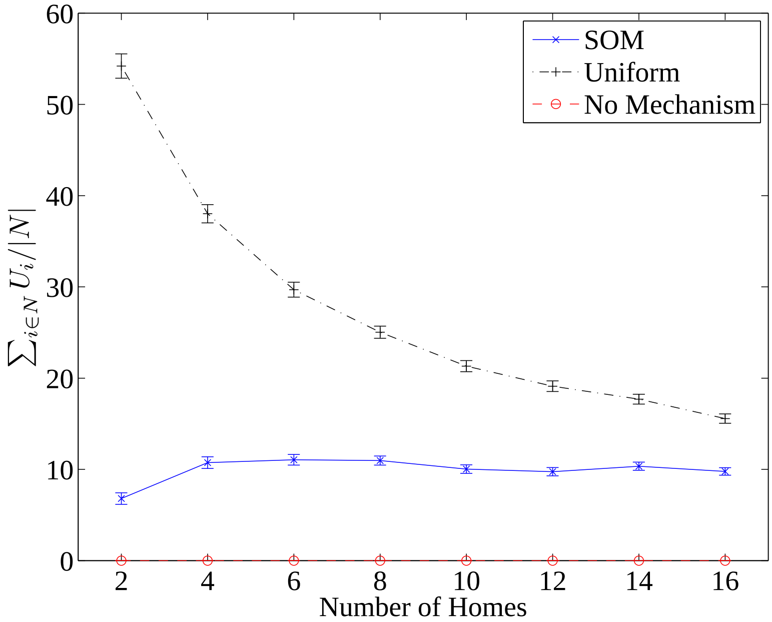

FIGURE 4.9: The additional utility received by the home agents under the uniform and sum of others' plus max mechanisms.

uniform mechanism here provides a greater amount of utility to the home agents than does the sum of others' plus max mechanism should not discourage the aggregator from implementing the latter mechanism. What is more important to the aggregator is the amount of utility it gains by implementing a given mechanism over using no mechanism. Empirical analysis of the aggregator's utility and risk is presented in the next section.

4.4.6 Risk and Utility to the Aggregator

Although Section 4.4.5 has shown that both the sum of others' plus max and the uniform mechanism are individually rational with respect to the home agents, individual rationality must also be shown from the perspective of the aggregator, which is itself a rational agent. This property of individual rationality for the aggregator corresponds to that of *budget balance*. The aggregator will make some savings through using the home agents' information, a portion of which it allocates as the budget for rewarding the home agents. Consequently, a *weakly* budget balanced mechanism will distribute *at most* all of the allocated budget and result in a positive utility for the aggregator as it has made some savings through using the mechanism in comparison to using only its priors. This is important because while for the purpose of this analysis it is assumed that the aggregator must somehow participate in the electricity markets, the aggregator is free to implement *whichever* mechanism gives it the greatest utility. Consequently, this section analyses the difference in utility the aggregator receives when using the uniform and sum of others' plus max mechanisms in comparison to using no mechanism.

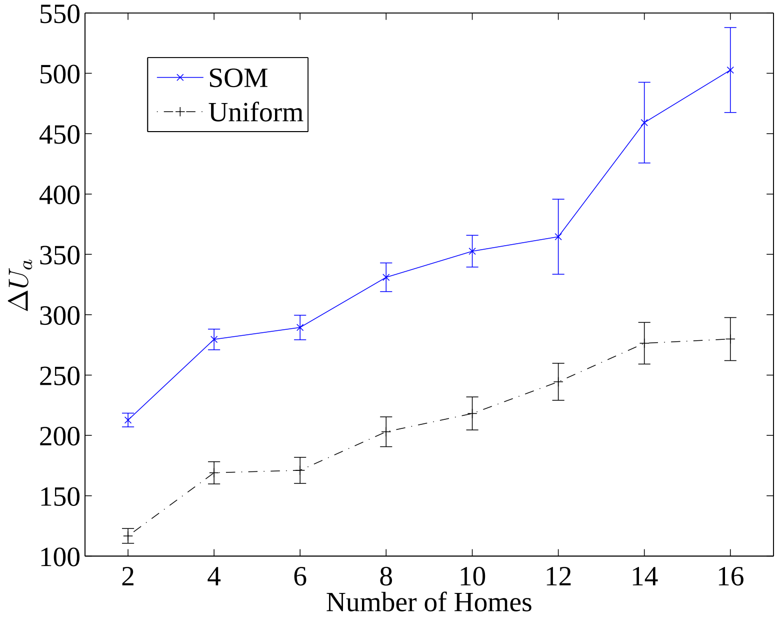

FIGURE 4.10: The average utility received by the aggregator under the uniform and sum of others' plus max mechanisms, with comparison to the aggregator's utility under no mechanism.

Figure 4.10 shows the additional utility gained by the aggregator when it uses the uniform or sum of others' plus max mechanisms compared to using no mechanism. As can be seen, the aggregator benefits from using either mechanism instead of no mechanism. This is fairly intuitive, as it can be seen from Section 4.4.4 that both of the mechanisms encourage agents to make precise predictions. Thus, the aggregator makes some amount of savings, only half of which it allocates to the reward of agents (as $\lambda = 0.5$, see Table 4.2 for an overview of the parameters). Furthermore, it can be seen that the sum of others' plus max mechanism provides greater additional utility to the aggregator than does the uniform mechanism. Again, taking into account the analysis of agents' precisions in Section 4.4.4, this is not surprising. However, even if all agents were to adopt the same strategies under both mechanisms, the aggregator would still obtain a greater utility from the sum of others' plus max mechanism. This is due to the fact that, whereas the uniform mechanism is *strictly* budget balanced (Theorem 4.1), the sum of others' plus max mechanism is *weakly* budget balanced (Theorem 4.7). Consequently, the uniform mechanism would distribute *all* of its allocated savings, whereas the sum of others' plus max mechanism would distribute *at most* all of its allocated savings.

In addition to absolute utility, it is interesting to analyse the *risk* to the aggregator. Risk is an important metric as it shows the possible worst case utilities for the aggregator. In particular, the metric used here is aggregator's utility *five percent value at risk*, which measures the upper bound on the *five percent worst* utilities that the aggregator receives. To calculate the five percent value at risk, the aggregator's utilities for all 1000 samples are sorted, and the 50^{th} lowest value is taken. Thus, in expectation, the aggregator

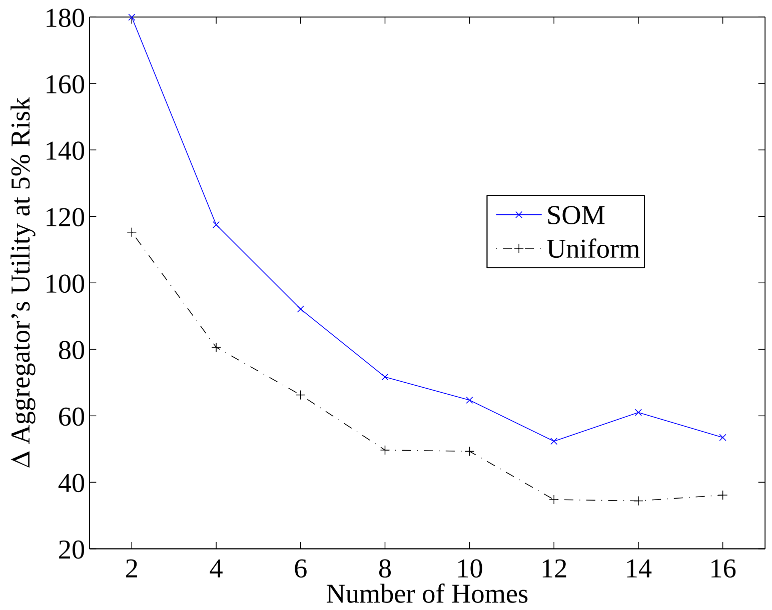

FIGURE 4.11: The value at five percent risk of the aggregator's utility for the uniform, sum of others' plus max, and no mechanism simulations.

should receive a utility greater than the five percent value at risk ninety-five percent of the time, and therefore an increased value at risk represents an *reduction* in risk.

Figure 4.11 shows the *difference* in the value of the aggregator's utility at five percent risk for the sum of others' plus max and the uniform mechanisms in comparison to that when simply using the aggregator's priors. Given this, it can be seen that the sum of others' plus max mechanism offers the aggregator a greater *reduction in risk* than the uniform mechanism when compared with using no mechanism. The gradient of Figure 4.11 is slightly misleading. The negative gradient does *not* represent the fact that the risk is increasing for the aggregator. This actually occurs because the risk is being reduced by the increased numbers (as was discussed in Section 4.4.4). Therefore, the five percent value at risk for no mechanism is increasing.

It is helpful here to discuss the source of the risk to the aggregator. Risk is an inherent feature of any mechanism in which a cost is incurred due to imprecision in the prediction of a future event. In terms of the aggregation scenario presented here, the aggregator's risk arises due to the fact that it will be penalised for over- or under-purchasing. Consequently, while the aggregator has an imprecise prediction – that is one that has finite precision – the aggregator will always expect to make *some* loss. This is the risk shown in Figure 4.11. Indeed, looking at the points representing the same numbers of agents, it can be seen that the mechanism that encourages the most precise predictions in Figure 4.8 (i.e. sum of others' plus max) is also the mechanism that gives the greatest reduction in the aggregator's risk. However, the aggregator will only be penalised if the amount it buys based upon the *aggregate* belief is different from the amount consumed. Furthermore, the distribution of the agents' consumptions is two-tailed (i.e. agents may

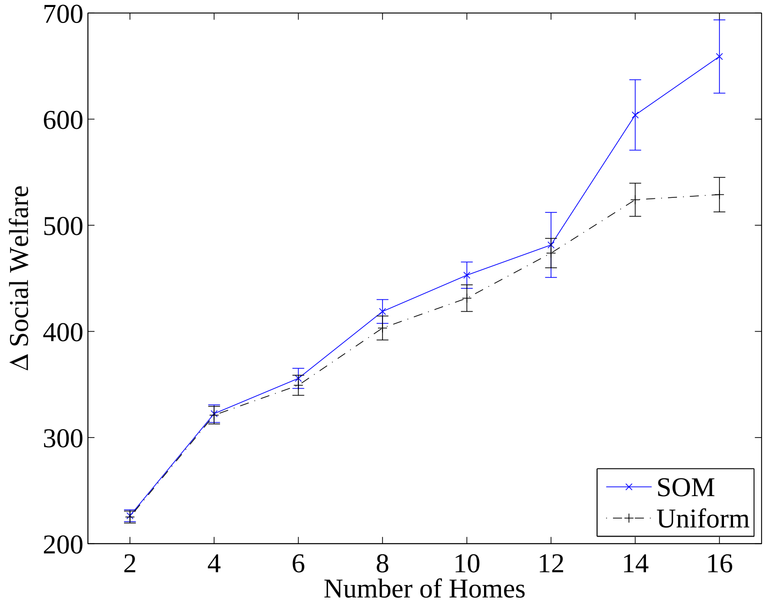

FIGURE 4.12: The additional utility received by the home agents under the uniform and sum of others' plus max mechanisms.

consume more or less than expected). Consequently, as the number of agents increases, so does the probability that two or more agents' errors may cancel one-another out. Thus, there is an inherent reduction in risk by increasing the number of agents over which the aggregator is aggregating.

Having discussed both the home agents' utility in Section 4.4.5 and the aggregator's utility here, it is now possible to discuss the *social welfare* that arises due to the equilibrium strategies in each of the mechanisms. This is discussed in the next section.

4.4.7 Social Welfare

The social welfare of each mechanism is taken to be the sum of the aggregator's and the home agents' utilities:

$$U_{sw} = U_a(\boldsymbol{x}, \boldsymbol{x}_a, \boldsymbol{\omega}) + \sum_{i \in N} U_i(\boldsymbol{x}, \boldsymbol{x}_a, \boldsymbol{\omega})$$

Looking at the definitions of U_a and U_i in Equations 3.11 and 3.13 respectively, it can be seen that this interpretation of the social welfare is in fact equivalent to the amount of money made by the aggregator in selling electricity to the home agents, minus the cost of purchasing that electricity from the markets, minus the cost incurred by the agents resulting from their chosen strategies.

Figure 4.12 shows the *change* in social welfare when using either the sum of others' plus max, or the uniform mechanisms compared to using no mechanism. It can be seen that for low numbers of homes, the two mechanisms are roughly equivalent. However, as the number of homes increases, the sum of others' plus max mechanism becomes more beneficial in terms of social welfare.

This is an interesting result because, as can be seen in Figure 4.8, the sum of others' plus max mechanism in fact encourages much more precise predictions from agents. Consequently, larger savings would be made by the sum of others' plus max mechanism. However, such precision is very costly to the home agents, and thus the home agents lose a large amount of utility in comparison the uniform mechanism (as can be seen in Figure 4.9).

4.5 Conclusions from this chapter

This chapter presented the *uniform* and *sum of others' plus max* mechanisms. The theoretical properties of both mechanisms were discussed along with proofs that the uniform mechanism is *strictly* budget balanced, and *Nash incentive compatible*, as well as proofs stating that the sum of others' plus max mechanism is *weakly* budget balanced, and *dominant strategy incentive compatible*. In addition to the theoretical properties of the mechanisms, this chapter furthermore analysed properties that emerged during the equilibrium points that developed as a result of the mechanisms. Due to the complexity of solving this analytically, computational simulation was used. It was shown that the sum of others' plus max mechanism encouraged home agents to make more precise predictions than under the uniform mechanism, which in turn increased social welfare and profit for the aggregator. However, it was shown that home agents gained a smaller amount of utility from the sum of others' plus max mechanism than they did from the uniform mechanism. Even so, both the uniform and sum of others' plus max mechanism were shown to be ex ante individually rational to both the home agents and to the aggregator. Furthermore, it was shown that the sum of others' plus max mechanism reduced the aggregator's risk.

Thus, under the scenario in which agents may manipulate their reports, it was shown that the sum of others' plus max mechanism improved upon the simpler uniform mechanism. The next chapter extends this analysis by relaxing the assumption that the agents' events are fixed. In it, the incentives provided by the mechanisms are analysed under the assumption that agents are able to make themselves harder or easier to predict.

Chapter 5

Mechanism Design for Scenarios with Load Manipulation

The previous chapter analysed the incentives provided by the uniform and sum of others' plus max mechanisms when the home agents are able to misreport their beliefs. However, it was assumed that the agents' associated events (i.e. ω_i for each agent i) are fixed. In reality, this assumption does not necessarily hold. Home agents and their respective home owners have control over their electricity consuming devices. Thus, they may also manipulate their consumptions in order to maximise their utility, in particular, they may manipulate their consumption such that the centre has a less accurate estimate of future consumptions compared to their own (as described in Section 3.3.2), or they may choose to waste electricity such that they always consume the amount they reported to the aggregator (see Section 3.3.3). This chapter investigates the incentives for agents to manipulate their loads. In doing so, Section 5.1 first analyses the incentive for agents to manipulate the *predictability* of their consumption, which the agents change by increasing or decreasing the variance of their loads. Since the agents manipulate their predictability by increasing or decreasing the variance of the loads they place on the grid, this form of manipulation is known as *load variance manipulation.* Following that, Section 5.2 analyses the incentive for agents to waste electricity such that they consume precisely what they predicted.

5.1 Load Variance Manipulation

To begin with, this chapter analyses the incentive for agents to make themselves easier or harder to predict by applying *load variance manipulation.* In essence, this scenario differs from that of the previous chapter in that agents are able to, at a cost, *manipulate the aggregator's priors.* From the perspective of a mechanism, it should be noted that incentivising this form of manipulation *does not* violate incentive compatibility. That

Symbol	Description	Value
—	Trials	1000
n	Number of agents	2–10
θ_{max}	Maximum scored precision	1500
λ	Fraction of savings to distribute	0.5
f	Forward market price	100
δ^b	System buy price delta	50
δ^s	System sell price delta	70
f_r	Retail price per unit of electricity that the agents must pay to the aggregator	150
$\alpha_{\theta,i}$	Agent i's precision cost coefficient	$i/100$
$\alpha_{\gamma,i}$	Agent i's load variance manipulation cost coefficient	$i/10$

TABLE 5.1: The values of parameters used by the simulation for Chapter 5.

is, a mechanism might well still be truth-revealing and therefore incentive compatible but it might at the same time encourage agents to manipulate their consumptions such that they are easier or harder to predict.

The remainder of this section is structured as follows: First, Section 5.1.1 discusses the configuration of additional parameters used in simulations within this section. Then, Section 5.1.3 discusses the theoretical incentives for agents to manipulate their consumptions. Next, Section 5.1.4 provides a discussion of the equilibrium strategies that are encountered under each of the mechanisms.

5.1.1 Simulation Setup

For the most part, the method used to obtain the empirical results within this chapter is the same as that from Chapter 4. Consequently, it is not discussed here and the reader is encouraged to see Section 4.4.1 for further information. All parameters, described in Table 5.1, are the same as those used in Chapter 4 with the exception of the numbers of agents, which has been reduced for computational reasons and an *additional* parameter denoting the agents' incurred costs for applying load variance manipulation, $\alpha_{\gamma,i}$, which is discussed below.

5.1.2 The Cost of Load Variance Manipulation

Similar to the precision costs of Chapter 4, differing costs of load variance manipulations among agents are represented by individual cost coefficients proportional to each agent's number. That is, each agent, i, has a load variance manipulation cost coefficient of $\alpha_{\gamma,i} = i/10$. Load variance manipulation is therefore more costly than creating precise predictions in that increasing the amount of load variance manipulation by one costs ten times as much as increasing the precision of an agent's report.

This seems to be a valid assumption as manipulation of load variance will conceivably require much more computation. This is because in order to achieve load variance manipulation, the home agent must schedule loads such that they are more or less predictable while not affecting the occupants of the house. In terms of demand response, this would require the home agent to differentiate between deferrable and non-deferrable loads. Moreover, the agent would need to ensure that the utility of the electricity used by the home-owners is not decreased. A step toward ensuring this is by guaranteeing that electricity-consuming tasks are still completed within deadlines imposed by the home-owners. For example, if a home owner were to put on their washing machine to be completed by the following morning. In this way, the cost of using more or less electricity will still be amortised by the utility of consuming that electricity. Furthermore, given this, analysis needs only to take into account the load variance manipulation costs, defined by Equation 3.9, and does not need to account for lost utility due to differing electrical consumption. This is because under load variance manipulation, the average consumption is unaffected and the home owners' loads are still all placed on the grid. It is just assumed that the home agent is able to shift loads such that they are harder or easier to predict while maintaining the home owner's utility from that load.

Given these costs and the additional utility gained by potentially increased rewards, agents will strategise over the amount of load variance manipulation they apply. The incentives for agents to manipulate their load variance is analysed in the next section.

5.1.3 Incentives to Manipulate

As previously stated, the aim of this chapter is to investigate the behaviours of agents under the uniform and sum of others' plus max mechanisms when agents are able to make themselves harder to predict. Given this, this section provides an analytical discussion of the incentives for agents to manipulate the predictability of their consumptions. This is done through the plotting and analysis of the agents' utility functions with respect to both their load variance manipulation γ_i and their precision of reports, θ_i. Through this, the behaviour of a single agent with respect to a set of opponents with a fixed strategy can be analysed. This is useful in order to understand the equilibria that are computed in Section 5.1.4.

It is assumed that even under no mechanism, the agents are able to manipulate their consumptions. Thus, this section begins by a discussion of the agents' incentives to perform load variance manipulation under the 'no mechanism' case. Following that, the agents' incentives under the uniform mechanism are discussed. Then, the section concludes with a discussion of the agents' incentives under the sum of others' plus max mechanism, before discussion turns to empirical analysis of these mechanisms in the next section.

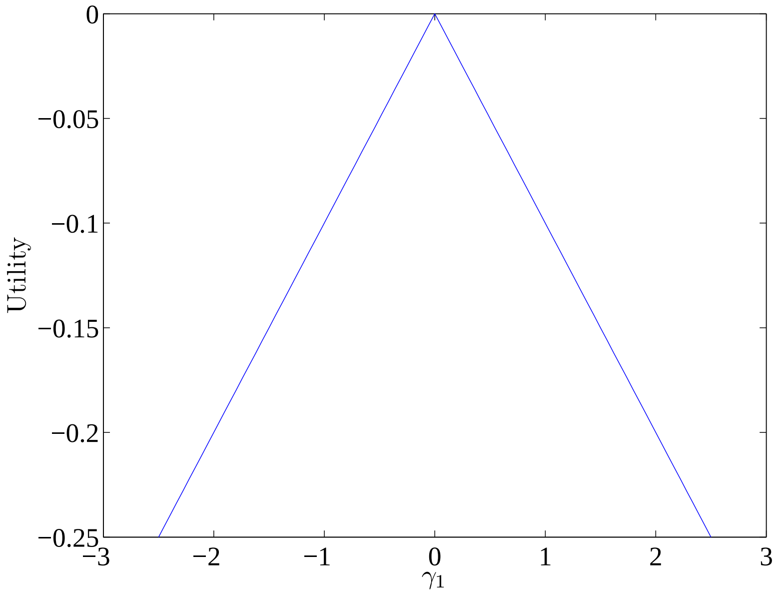

FIGURE 5.1: The expected utility gained by an agent by performing varying amounts of load variance manipulation, γ_1, in the 'no mechanism' scenario.

5.1.3.1 No Mechanism

The incentive for agents to manipulate the predictability of their loads comes from the fact that the *savings*, which determines the budget for rewarding agents, is dependent on the *improvement in precision* between the agents' beliefs and the *aggregator's priors*. That is, if the agents can make the aggregator's priors less precise while maintaining the precision of their own belief, the agents will increase the budget from which they are rewarded, thereby increasing their reward.

Of course, in the 'no mechanism' scenario, agents are not rewarded at all. Thus, manipulating their predictability makes no difference to their utility other than a decrease due to the cost incurred through manipulation. This can be seen in Figure 5.1, which shows agent 1's expected utility from deviating from its equilibrium strategy (found using iterated best response, as discussed in Section 3.6.1). It can be seen that any manipulation of the agent's predictability causes a loss to that agent, as was discussed previously.

However, this is not the case for all mechanisms. For example, the uniform mechanism, which is discussed next, encourages agents to make themselves harder to predict.

5.1.3.2 Uniform

Unlike the 'no mechanism' scenario, under the uniform mechanism, agents are paid based upon the *improvement* they make to the aggregate belief when compared to the aggregator's priors. The uniform mechanism, as described by Equation 4.1, defines

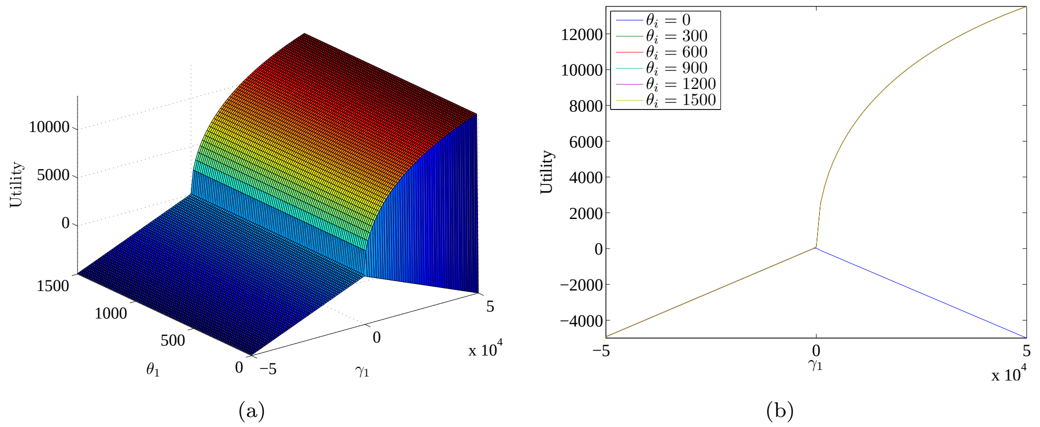

FIGURE 5.2: The expected utility gained under the uniform mechanism by agent 1 when reporting a precision of θ_1, while applying γ_1 load variance manipulation.

payments which are *directly proportional* to the savings. Therefore, if agents are able to increase the savings made by the aggregator, they will receive a higher reward.

Under this model, there are two methods of achieving this increase in savings: increasing the precision of the agents' reports, or *decreasing* the precision of the *aggregator's priors*. However, as discussed in Section 3.2, Section 3.3.2 and Section 5.1.2, there are various costs involved in both of these methods. Therefore, an agent must strategise in order to choose a precision-load-variance-manipulation pair that maximises its utility.

The plots in Figure 5.2 show agent 1's utility function for applying varying degrees of load variance manipulation, γ_i. It can be seen that the utility gained by applying positive load variance manipulation far outweighs the costs incurred. Thus, agents choose to apply a large amount of manipulation, thereby making the aggregator increasingly imprecise. Of course, the exact amount of load variance manipulation applied by the agents is dependent upon their costs. Thus, for agents with high costs this effect will be far less pronounced, with the optimal γ_i approaching zero as $\alpha_{\gamma,i}$ increases.

5.1.3.3 Sum of Others' plus Max

While the uniform mechanism, as discussed previously, encourages agents to make themselves *harder* to predict, this is *not* the case for the sum of others' plus max mechanism. Figure 5.3 shows agent 1's expected utility for applying load variance manipulation, γ_1 while reporting information of precision, θ_1. It can be seen that under sum of others' plus max, with the parameters described in Table 5.1, agents have an incentive to manipulate their load variance such that they are *easier* to predict.

This comes about due to the fact that the aggregator's prior for the agent being rewarded, i, is actually substituted in place of agent i's. Recall from Equation 4.8, that the savings

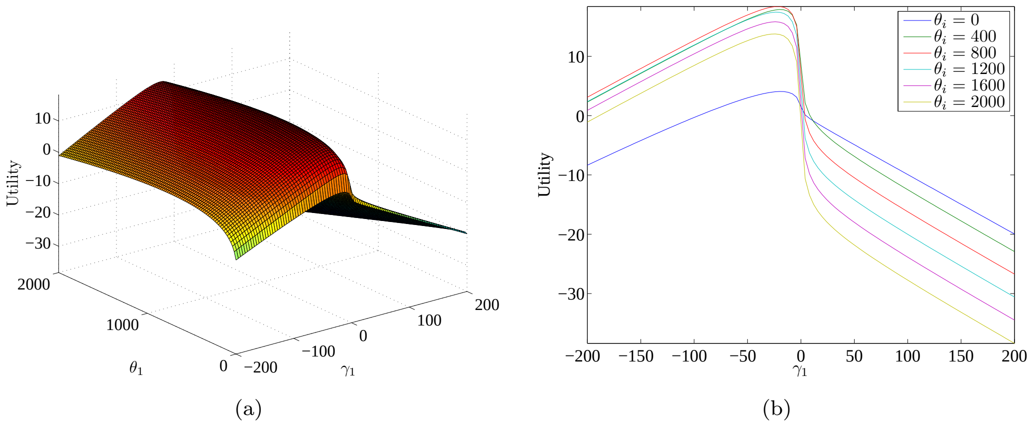

FIGURE 5.3: The expected utility gained under the sum of others' plus max mechanism by agent 1 when reporting a precision of θ_1, while applying γ_1 load variance manipulation.

used to determine the budget are in fact:

$$\lambda \cdot \Delta\left(\boldsymbol{x}_{-i} \cup \{x_{a,i}\}, \boldsymbol{x}_a, \boldsymbol{\omega}\right)$$

Thus, if an agent applies positive load variance manipulation, it will in fact *decrease* the precision of the aggregate belief used to calculate the savings made by the aggregator (i.e. $\boldsymbol{x}_{-i} \cup \{x_{a,i}\}$). However, because its own report is substituted by the aggregator's prior, by making itself *easier* to predict, the agent also *increases* the precision of the aggregate belief. In this manner, an agent making itself easier to predict will increase the budget from which it is rewarded and as a consequence, with all other variables remaining the same, it will increase its reward. This behaviour is of great advantage to the grid as it reduces the amount of standby generation that is required and allows planners to make more optimal decisions regarding the use of grid hardware.

While it is known now that agents will make themselves harder to predict in the uniform mechanism and easier to predict under sum of others' plus max, it is not known precisely how agents will behave with regards to the amount of load variance manipulation and the precision of the agents' reports. Moreover, it is not known how these behaviours will affect the precise utility of the aggregator and the home agents. Consequently, the next section is devoted to empirical analysis of the mechanisms under the scenario with load variance manipulation.

5.1.4 Empirical Evaluation of the Mechanisms

Having looked at the analytical properties of the mechanisms in the previous section, and having discussed the fact that agents will make themselves easier to predict under the uniform mechanism and harder to predict under sum of others' plus max, this section discusses the empirical analysis of the equilibria that arise under these mechanisms.

Specifically, this section begins with a brief discussion regarding the calculation of savings and how they no longer necessarily relate to the benefit the agents provide to the aggregator. Next, the load variance manipulation applied by the agents and its effect on the aggregator's priors is discussed. Following that, the precision of reports from the home agents is discussed. Then, there is a discussion of the utility of the home agents under the mechanisms, as well as the utility and risk to the aggregator. Finally, the section ends with the analysis of the social welfare achieved by the mechanisms.

5.1.4.1 A Discussion regarding Savings

The use of savings within the mechanisms presented in this thesis was first proposed due to the fact that savings seemed to clearly indicate the benefit an agent made to the aggregator. That is, if the cost of electricity is ten when the aggregator uses its prior information but the cost of the same electricity is five when using the agents' information, the agents obviously caused a benefit to the aggregator of value five. However, this is not true for the scenario presented in this chapter. Agents are able to increase the *perceived* savings by making themselves harder to predict. Thus, a large amount of savings can come about not only through actions which benefit the aggregator (i.e. providing high precision estimates), but also actions which damage the aggregator (making the agents harder to predict). Furthermore, an agent making itself easier to predict in fact *reduces* the cost of electricity for the aggregator when using its priors. In turn, this *reduces* the perceived savings made by the aggregator.

The reason for this is as follows: When deciding the optimal amount to buy, it is assumed that the aggregator has no information regarding the *maximum consumption* of a given home. This results in an unbounded amount of savings as $\gamma_i \to \infty$. When purchasing electricity, the aggregator will in fact purchase an amount of $\mu + \varepsilon$ where ε is some offset that minimises the expected cost of electricity. Due to the fact that the system sell price delta is greater than the system buy price delta, this results in ε being positive. That is, the aggregator chooses to purchase excess electricity as it knows that buying more electricity in the balancing market is more costly than the loss made by buying excess units in the forward market and selling them back to the grid in the balancing market. However, with no information regarding the maximum consumption of a home, the aggregator will potentially buy an unbounded amount of electricity for the homes. Consequently, as the aggregator's priors tend to zero precision, so the amount to buy increases. This results in large amounts of savings when compared to the use of the home agents' reports, which results in the aggregator purchasing an amount of electricity much closer to the realised event.

It is this peculiarity that determines how agents behave with regards to their predictability. Indeed, it has already been discussed how this fact encourages agents under the

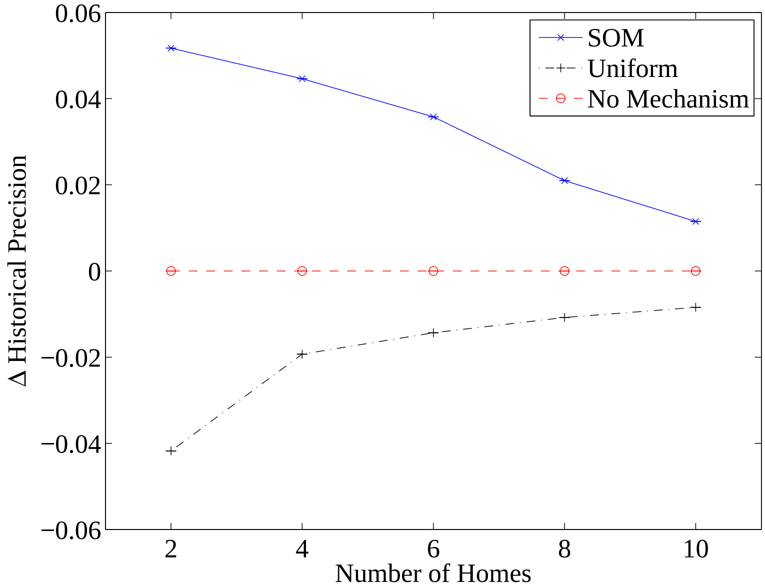

FIGURE 5.4: The difference in historical load precision between agents' *unmanipulated* consumptions and their historical loads *after* manipulation.

uniform mechanism to make themselves harder to predict. This is important to remember when interpreting the upcoming empirical results, as it shows how under certain conditions, the aggregator is able to make a *loss* in comparison to the 'no mechanism' scenario – something which would not be possible with budget balanced mechanisms and no load variance manipulation.

Discussion of the empirical results begins in the next section with a discussion of how agents exploit this subtlety in the *perceived* savings through load variance manipulation.

5.1.4.2 Load Variance Manipulation

Under the model described in Section 3.2 and Section 3.3.2, agents may apply a form of manipulation known as *load variance manipulation*. As defined in Equation 3.8, applying positive load variance manipulation, (i.e. $\gamma_i > 0$) results in agents making themselves *harder* to predict – they increase the variance of their loads. Conversely, agents applying negative load variance manipulation make themselves *easier* to predict. This in turn affects the aggregator's priors in the same manner.

As was discussed in Section 5.1.3, the uniform mechanism encourages agents to make themselves harder to predict as doing so increases the budget from which they are rewarded, and thus increases their reward. However, as has been discussed, the sum of others' plus max mechanism encourages agents to make themselves *more* predictable. Figure 5.4 shows the *change* in the aggregate historical precision due to load variance manipulation. Positive numbers represent agents making themselves more predictable whereas negative numbers indicate that agents are making themselves harder to predict.

Therefore, it can be seen that empirical analysis corroborates the theoretical discussion of Section 5.1.3 in that under the uniform mechanism agents make themselves harder to predict, whereas under sum of others' plus max, agents make themselves easier to predict.

A key advantage to the behaviour of the home agents under sum of others' plus max is that, reducing agents' load variance implies less volatility of demand within the grid. This volatility, as was discussed in Section 2.2, is the cause of a large amount of expenditure and carbon output in energy production. Consequently, while simply making accurate predictions allows for a reduction of 'spinning reserve' energy, a lower volatility of demand allows for a lower peak capacity to be supported within the grid. This is important as for the majority of the time, the full capacity of the grid is unused. Thus this capability is being maintained only for use a small amount of the time. Lowering this peak capacity will reduce overheads in maintaining this capability, and thus will further reduce electrical cost, and may even reduce the carbon output of the grid.

However, while negative load variance manipulation is advantageous to the grid, the key to ensuring a low price of electricity for the aggregator is to reduce its exposure to the balancing markets. This occurs through using precise predictions from the home agents themselves, and the precision of those beliefs under equilibrium are discussed next.

5.1.4.3 Precision of Reports

Figure 5.5 shows the precision of the aggregate belief used by the aggregator under the uniform and sum of others' plus max mechanisms as well as under no mechanism. It is important to note that due to the fact that the aggregator does have prior information, the red, 'no mechanism' line is in fact greater than zero. However, the precision of the aggregate belief under 'no mechanism' is dwarfed by the aggregate precision under the uniform and sum of others' plus max mechanisms.

The results from Figure 5.5 show that for the most part, the sum of others' plus max mechanism outperforms the uniform mechanism in terms of the aggregate precision. However, for the two agent scenario, the uniform mechanism vastly outperforms sum of others' plus max.

At first glance this seems to be an interesting result as under the uniform mechanism, agents seem to be placing their efforts into opposing forces. That is, on the one hand, the agents are making themselves harder to predict, whilst they are themselves providing precise estimates. In fact, this behaviour arises due to the agents' incentive to *increase the perceived savings* made by the aggregator. By making themselves harder to predict, the home agents make electricity much more expensive for the aggregator to purchase using only its priors. Then, by themselves giving a high precision estimate, they make

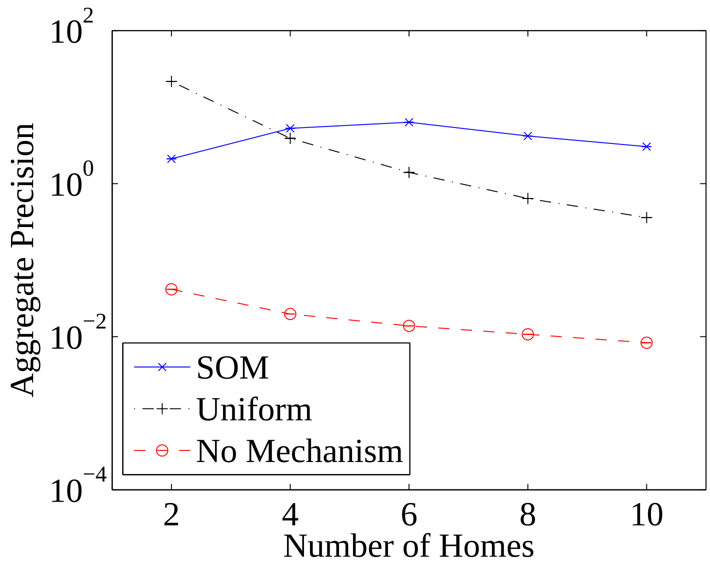

Figure 5.5: The precision of the aggregate belief used by the aggregator in order to purchase electricity under the uniform and sum of others' plus max mechanisms as well as under no mechanism.

the amount of money *actually* spent by the aggregator using their beliefs smaller. Thus, the savings are increased, and consequently so are the agents' utilities.

Under sum of others' plus max, the agents choose to make themselves easier to predict. However, they still provide a precise estimate of their future demand, albeit less precise than the uniform mechanism for numbers of agents less than four. The agents' behaviour regarding predictability has been discussed already in Section 5.1.3. However, the agents' incentive to provide estimates is not dependent on increasing the savings made by the aggregator. Instead, it is dependent on maximising the score the agent will receive given the cost they incur due to their precision. Furthermore, as defined in Equation 3.4, as agents increase their predictability, the cost of producing an estimate of a given precision reduces. Consequently, the precision of reports generated by the home agents remains relatively high, even when savings are being divided among increasing numbers of agents. It can be seen that this is not the case for the uniform mechanism due to rewards being diluted among the agents.

Even so, as is discussed in the next section, the agents' strategies under the uniform mechanism can be incredibly profitable for them. Although, as is discussed afterwards, this is to the detriment of the aggregator.

5.1.4.4 Utility to the Home Agents

It has already been discussed that under the uniform mechanism, home agents seek to maximise the savings made by the aggregator by making themselves harder to predict. It can be seen from Figure 5.6 that this is incredibly profitable for the home agents, who

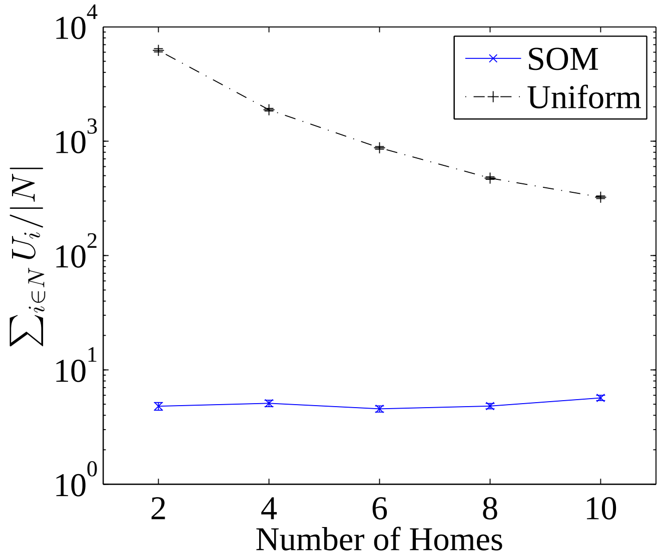

FIGURE 5.6: A comparison of the increase in utility to the home agents by making themselves harder to predict under the sum of others' plus max mechanism and the uniform mechanism over using no mechanism.

receive a utility of up to the order 10^3. Of course, as was the case in the no manipulation scenario of Chapter 4, the utility per agent decreases as the number of agents increases.

It should be noted that the sum of others' plus max mechanism *does* result in positive utility for the home agents. However, the utility gained by the home agents is much lower, in the order of 10.

Initially, the increase in utility received by the home agents in the uniform mechanism might appear to be a good thing. However, this increase in utility does not simply come from agents supplying improved predictions. Indeed, it can be seen from Figure 5.5 that for the most part, the precision of the home agents' reports under the uniform mechanism is far below the precision of the home agents under the sum of others' plus max mechanism. In fact, this extra utility comes from the agents' load variance manipulation, *damaging* the aggregator's priors by making themselves harder to predict. This incentive is precisely the opposite to the one desired. Furthermore, the next section shows that the home agents' behaviour under the uniform mechanism causes a huge loss in utility to the aggregator compared to if the aggregator were to simply have used its priors.

5.1.4.5 Risk and Utility to the Aggregator

While the previous section showed that under the uniform mechanism the utility to the home agents is high, this section will show that this comes at the expense of the aggregator. It is shown that the aggregator fares much worse under the uniform mechanism

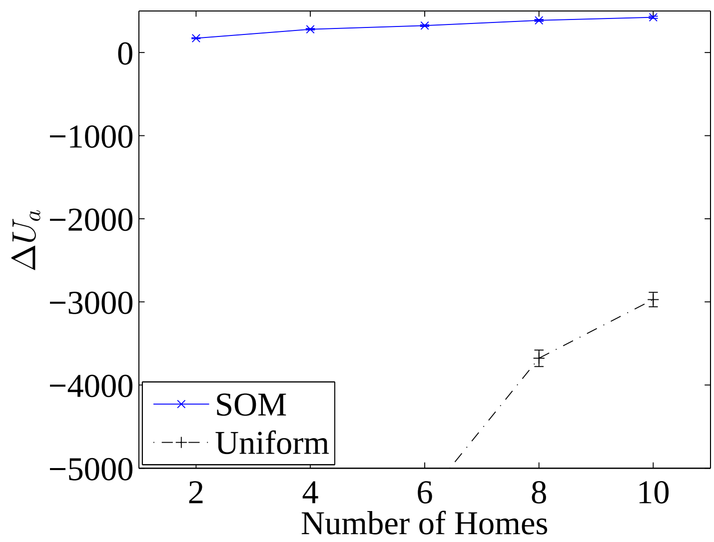

FIGURE 5.7: The change in the aggregator's utility from using the uniform or sum of others' plus max mechanisms rather than using no mechanism.

than under the sum of others' plus max mechanism, effectively writing off the uniform mechanism as a viable mechanism for use in scenarios with load variance manipulation.

Figure 5.7 shows the difference in utility received by the aggregator when using the sum of others' plus max and the uniform mechanisms compared to if it were to use no mechanism. That is, the difference in using the agents' information with load variance manipulation in comparison to using the aggregator's priors with no load variance manipulation (as was discussed in Section 5.1.3). It can be seen that the aggregator loses a large amount of utility under the uniform mechanism. As the number of agents increases, this utility increases. However, while there is positive load variance manipulation within a mechanism, the aggregator will always lose utility in comparison to using no mechanism. This is because the agents are actively crippling the aggregator's priors. Consequently, using only the aggregator's priors, the cost of electricity is much greater and, as the agents' rewards are proportional to the *difference* in the cost of purchasing electricity using the priors and the agents' reports, the aggregator pays out a larger amount of money to the agents.

Figure 5.8 shows the worst case difference in utility (at five percent risk) between the aggregator using no mechanism and the aggregator using the uniform and sum of others' plus max mechanisms. It can be seen that under the uniform mechanism, this difference in worst-case utility is large, whereas under sum of others' plus max, the worst-case difference is very small (approaching zero for larger numbers of agents). Regardless of the mechanism used, there will always be individual instances in which the aggregator's utility is negative due to the fact that while the precision of the reports used to purchase electricity is finite, there will always be rare occasions in which the prior information held by the aggregator was more *accurate* (i.e. closer to the realised consumption) than

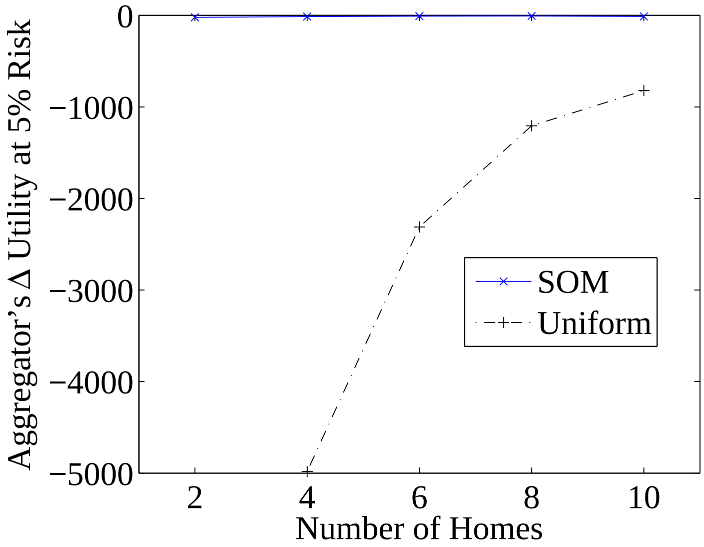

FIGURE 5.8: The change in the aggregator's utility at five percent risk when using the uniform and sum of otters' plus max mechanisms rather than using no mechanism.

the agents reports. Therefore, this shows that the sum of others' plus max mechanism is of far less risk to the aggregator than the uniform mechanism; Even in outlying events in which the centre's prior information is more accurate than the agents' reports (which becomes increasingly less likely as the precision of the agents' reports increases) the aggregator's loss is greatly reduced.

Clearly then, a rational aggregator would choose to implement the sum of others' plus max mechanism as it guarantees the aggregator not to make a loss in expectation. This occurs due to the fact that agents' incentives are towards making themselves easier to predict, reducing the cost to the aggregator. Furthermore, as the home agents' utility and the aggregator's expected utilities are always positive, using sum of others' plus max also increases social welfare, as is discussed next.

5.1.4.6 Social Welfare

It has already been discussed how the uniform mechanism is detrimental to the aggregator's utility whereas the sum of others' plus max mechanism improves it. However, it has also been discussed that performance in terms of the home agents' utility is greatly improved in the uniform mechanism compared to the sum of others plus max. Thus it is interesting to analyse the social welfare afforded by the two mechanisms under equilibrium, a plot of which is shown in Figure 5.9.

It can be seen from Figure 5.9 that the sum of others' plus max mechanism improves the social welfare compared to the uniform mechanism. Moreover, Figure 5.9 shows the uniform mechanism to actually be detrimental to social welfare for numbers of agents less than 8. Thus, the sum of others' plus max mechanism provides a marked benefit to

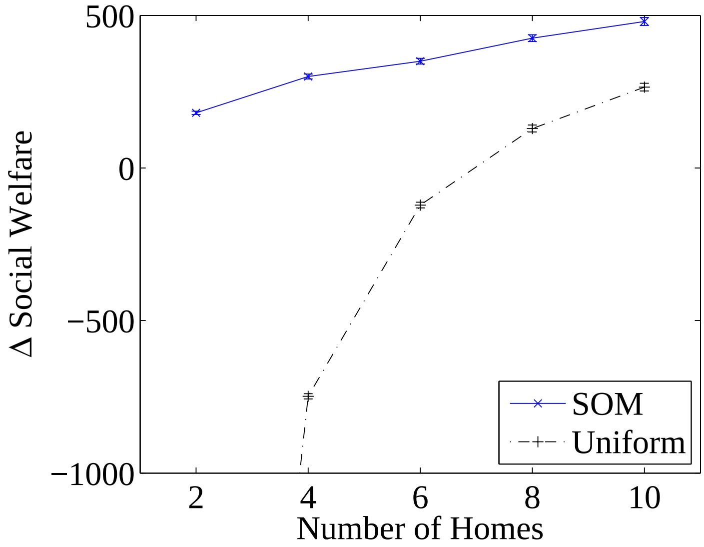

Figure 5.9: The change in social welfare from using the uniform and sum of others' plus max mechanisms over using no mechanism.

the entire community of agents, although that difference can be seen to be reduced as the numbers of agents increases.

Of course, this discussion is fairly academic as a rational profit driven aggregator would implement the sum of others' plus max mechanism regardless. However, this discussion shows that even if an aggregator wanted to maximise social welfare regardless of its own profit, it would still choose the sum of others' plus max mechanism over the uniform mechanism.

5.2 Incentives to Waste

So far in this chapter, discussion regarding the manipulation of loads has been restricted to agents making themselves harder or easier to predict. Of course, as well as manipulating their predictability, an agent that is able to manipulate its realised consumption might also be able to *waste* electricity in order to ensure its consumption matches its prediction. This form of manipulation is formalised in Chapter 3, specifically Section 3.3.3.

The remainder of this section continues as follows: Firstly, Section 5.2.1 discusses the theory regarding why the incentive to waste exists. Next, Section 5.2.2 discusses the incentive to waste electricity under the uniform mechanism. Then Section 5.2.3 discusses the incentive to waste electricity under the sum of others' plus max mechanism. Following that, Section 5.2.4 discusses how this incentive to waste may be removed from the mechanisms.

5.2.1 The Origin of Wasting

Work thus far in this thesis has made the implicit assumption that once agents report their predictions to the aggregator they receive no further information. However, in reality this is unlikely to be the case. In a real implementation, the last point at which an agent may supply information to the aggregator is likely to be just before *gate closure*. In this context, gate closure is the time at which transactions for a specific time period may no longer be carried out within the forward markets. In many systems, gate closure occurs an hour before the time for which electricity is being purchased. That is, if electricity is being purchased for the time 15:00–16:00, gate closure for that period is 14:00 and any further alterations must be performed in the balancing markets.

It has been assumed so far within this thesis and to a large extent it has been assumed within the literature, that agents gain no new information after they send their reports (or at least, that agents are unable to manipulate the event which they are predicting). However, assuming that the home agents within this thesis are able to constantly gather information, it is conceivable that between 14:00 and the 16:00 (to use the example above), the home agents gather information that leads them to believe that they will no longer consume as they reported.

Furthermore, coupled with the assumption that agents are able to manipulate the amount they consume to match their beliefs, it is possible for some particularly undesirable incentives to arise. On the one hand, if agents were to learn that they were to over consume, they might be encouraged to *reduce* their consumption. However, on the other hand, if agents were to learn that they were to *under consume*, the agents might well be incentivised to *waste* electricity in order to match their reports. It is this *wasting* with which the discussion in this sections is particularly concerned.

Given agent i's utility function, $U_i\left(\boldsymbol{x}, \boldsymbol{x}_a, \boldsymbol{\omega}\right)$, an agent will *waste* electricity if:

$$\frac{\mathrm{d}\, U_i\left(\boldsymbol{x}, \boldsymbol{x}_a, \boldsymbol{\omega}\right)}{\mathrm{d}\, \omega_i} > 0 \tag{5.1}$$

Thus, in order to prove that agents have an incentive to waste under a particular mechanism, it is simply necessary to show that Equation 5.1 holds given the rewards received by the agents under that mechanism. This analysis is performed for the uniform mechanism in the next section.

5.2.2 Wasting under the Uniform Mechanism

As was mentioned previously, agents will choose to waste electricity if they expect to gain utility in so doing. Thus, in order to determine whether or not agents have incentive to waste under the uniform mechanism, the first derivative of the agents' utility function with respect to their consumption ω_i must be taken. Formally, it must be shown that there is no positive gradient in the agents' utility functions. That is, for there to be no incentive for agent i to waste, the following must hold:

$$\frac{\mathrm{d}}{\mathrm{d}\omega_i}\left[P_i\left(\boldsymbol{x},\boldsymbol{x}_a,\boldsymbol{\omega}\right)+P_\omega\left(\omega_i,\ddot{\omega}_i\right)-f_r\cdot\omega_i-C_\theta\left(\theta_i,\alpha_{\theta,i}\right)-C_\gamma\left(\gamma_i,\alpha_{\gamma,i}\right)\right]<0 \tag{5.2}$$

A number of these terms are in no way dependent upon ω_i. Therefore, the problem can be simplified to:

$$\frac{\mathrm{d}}{\mathrm{d}\omega_i}\left[P_i\left(\boldsymbol{x},\boldsymbol{x}_a,\boldsymbol{\omega}\right)+P_\omega\left(\omega_i,\ddot{\omega}_i\right)-f_r\cdot\omega_i\right]<0 \tag{5.3}$$

Furthermore, this discussion is primarily concerned with the incentive for agents to *waste* their electricity. That is, once the agents have consumed their required amount, $\ddot{\omega}_i$ where $\ddot{\omega}_i<\mathbb{E}\left(\hat{x}_i\right)$, the agents consume yet more electricity such that in total they consume ω_i, with $\ddot{\omega}_i<\omega_i\leq\mathbb{E}\left(\hat{x}_i\right)$.

Recalling from Equation 3.12 that while an agent's consumption, ω_i, is less than or equal to its required consumption, $\ddot{\omega}_i$, the cost of purchasing electricity is amortised by the utility gained through its consumption. Further, recall that for every unit of electricity purchased greater than the amount actually required, the agent gains *no* extra utility but still incurs the cost of purchasing said electricity. Therefore, by assuming that $\omega_i>\ddot{\omega}_i$, Equation 5.3 can be further simplified to the form:

$$\frac{\mathrm{d}}{\mathrm{d}\omega_i}\left[P_i\left(\boldsymbol{x},\boldsymbol{x}_a,\boldsymbol{\omega}\right)-f_r\cdot\left(\omega_i-\ddot{\omega}_i\right)\right]<0 \tag{5.4}$$

Thus, in order for there to be an incentive to manipulate, the following must hold:

$$\frac{\mathrm{d}}{\mathrm{d}\omega_i}P_i\left(\boldsymbol{x},\boldsymbol{x}_a,\boldsymbol{\omega}\right)>f_r \tag{5.5}$$

expanding P_i to give

$$\frac{\mathrm{d}}{\mathrm{d}\omega_i}n^{-1}\cdot\lambda\cdot\Delta\left(\boldsymbol{x}_a,\boldsymbol{x},\boldsymbol{\omega}\right)>f_r \tag{5.6}$$

As can be seen in Equation 5.6, the only term on the left hand side of the inequality that is dependent upon the agents consumption, ω_i, is the savings function, which is defined in Equation 3.5 as the difference in the cost of electricity between when the aggregator uses its own priors and when it uses the agents' reports. However, for the purpose of this section, only the case in which an agent consumes *less* than the expected value of its report is of importance. Furthermore, because the aggregator will always purchase an amount *greater* than the expected value of the reported distribution (due to over-consuming being penalised more heavily than under-consuming), $\omega_i+\omega_{-i}<\chi\left(\boldsymbol{x}\right)$.

Therefore, the first derivative of the savings with respect to ω_i is as follows:

$$\begin{aligned}\frac{\mathrm{d}}{\mathrm{d}\omega_i}\Delta\left(\boldsymbol{x}_a,\boldsymbol{x},\boldsymbol{\omega}\right) &= \frac{\mathrm{d}}{\mathrm{d}\omega_i}\overbrace{\left[f\cdot\chi\left(\boldsymbol{x}_a\right)+\left(\omega_i+\omega_{-i}-\chi\left(\boldsymbol{x}_a\right)\right)\left(f-\delta^b\right)\right]}^{\text{Cost with priors}}\\ &\quad-\frac{\mathrm{d}}{\mathrm{d}\omega_i}\overbrace{\left[f\cdot\chi\left(\boldsymbol{x}\right)+\left(\omega_i+\omega_{-i}-\chi\left(\boldsymbol{x}\right)\right)\left(f-\delta^b\right)\right]}^{\text{Cost with agents' reports}}\\ &= \frac{\mathrm{d}}{\mathrm{d}\omega_i}\left[f\cdot\chi\left(\boldsymbol{x}_a\right)-\chi\left(\boldsymbol{x}_a\right)\left(f-\delta^b\right)\right]-\frac{\mathrm{d}}{\mathrm{d}\omega_i}\left[f\cdot\chi\left(\boldsymbol{x}\right)-\chi\left(\boldsymbol{x}\right)\left(f-\delta^b\right)\right]\end{aligned}$$

all terms of which are independent of ω_i. Therefore:

$$\frac{\mathrm{d}}{\mathrm{d}\omega_i}\Delta\left(\boldsymbol{x}_a,\boldsymbol{x},\boldsymbol{\omega}\right)=0$$

Thus, the savings made by the aggregator are constant with respect to the actual amount consumed, assuming that the amount to buy with both the priors and the agents' reports are greater than the amount that is to be consumed by the agents without manipulation. Consequently, $\mathrm{d}U_i/\mathrm{d}\omega_i < f_r$, and therefore under the uniform mechanism, even when agents are certain that they are not going to consume as they predicted, they do not have incentive to waste electricity.

5.2.3 Wasting under the Sum of Others' plus Max Mechanism

The sum of others' plus max mechanism is more complicated to prove or disprove wasting. In fact, using the inequality from Equation 5.5 and expanding the agent's reward function, the sum of others' plus max mechanism gives an incentive to waste if:

$$\frac{\mathrm{d}}{\mathrm{d}\omega_i}S\left(x_i,\omega_i\right)\cdot\psi > f_r \tag{5.7}$$

where

$$\psi=\frac{\lambda\cdot\Delta\left(\boldsymbol{x}_{-i}\cup\left\{x_{a,i}\right\},\boldsymbol{x}_a,\boldsymbol{\omega}\right)}{S_{max}+\sum_{j\in N\setminus\{i\}}S\left(x_j,\omega_j\right)} \tag{5.8}$$

is constant with respect to ω_i. Thus, the incentive to waste within the sum of others' plus max mechanism is dependent upon the actual scoring rule used. In addition, the incentive to waste is dependent upon the actual *distribution* used and agent i's reported precision.

First of all, to prove whether or not agents are able to gain utility by wasting, it can be seen that when ψ is constant, the upper bound in the additional utility an agent gets by increasing its consumption by $\mathrm{d}\omega_i$ is dependent upon the upper bound of the first derivative of the scoring rule. Figure 5.10 shows the first derivative of the spherical scoring rule with respect to ω_i. The interesting feature of this plot is that $\mathrm{d}S/\mathrm{d}\omega_i$

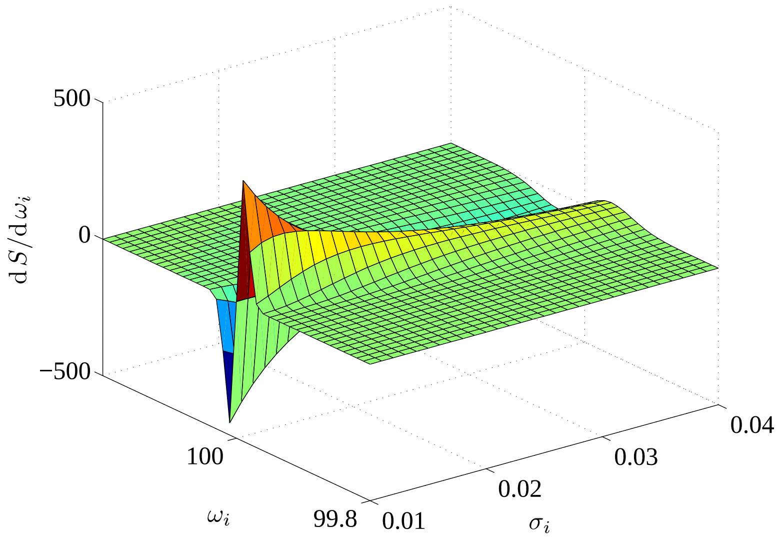

FIGURE 5.10: The first derivative of the spherical scoring rule with respect to ω_i at ω_i when the agent reports a belief of *standard deviation* σ_i.

increases as σ_i approaches zero. Furthermore, as $\sigma_i \to 0$ and $\omega_i \to \mathbb{E}(x_i)$, so $\mathrm{d}S/\mathrm{d}\omega_i \to \infty$. Of course, the actual maximum of the first derivative under sum of others' plus max is determined by the maximum precision, θ_{max}. Consequently, the maximum incentive for an agent to waste at any consumption ω_i under sum of others' plus max is:

$$\max\left(\frac{\mathrm{d}}{\mathrm{d}\omega_i}\psi \cdot S(x_i, \omega_i)\right) = \frac{\mathrm{d}}{\mathrm{d}\,\omega_i}\psi \cdot S(x_i, \omega_i^*) \tag{5.9}$$

where x_i is of maximum precision (i.e. $x_i = \langle \mu_i, \theta_{max} \rangle$) and ω_i^* is to be defined in the upcoming paragraphs. This will be different for each scoring rule used. Consequently, the aim of this discussion is just to provide instruction on how to determine whether an agent might have incentive to waste, and if so, how to remove that incentive. Within this discussion, in line with the rest of this thesis, this section assumes the *spherical* rule from Table 4.1. Furthermore, the derivative of the scoring rule is itself dependent upon the type of distribution used. Therefore, for mathematical convenience, Gaussian distributions are used throughout the rest of the analysis in this section. Furthermore, as is discussed in Appendix B, given certain constraints on the expected value and variance of a log-normal distribution, it can be approximated using a Gaussian distribution. Given these assumptions, the first derivative of agent i's score is:

$$\frac{\mathrm{d}}{\mathrm{d}\,\omega_i}S(x_i; \omega_i) = \frac{1}{\sqrt{\int_{-\infty}^{\infty}\mathcal{N}(\omega'; \mu_i, \theta_i)^2\,\mathrm{d}\,\omega'}} \cdot \frac{\mathrm{d}}{\mathrm{d}\,\omega_i}\mathcal{N}(\omega_i; \mu_i, \theta_i) \tag{5.10}$$

Thus to find the consumption at which an agent has a maximum incentive to waste electricity, the second derivative of the probability density function of the agent's report must be solved for ω_i. The solution for this gives:

$$\omega_i^* = \frac{\pm 1 + \mu_i \cdot \sqrt{\theta_i}}{\sqrt{\theta_i}} \tag{5.11}$$

Furthermore, as this section is concerned with the peak incentive to *waste*, solutions greater than μ_i can be removed. These solutions represent situations in which an agent learns that it will consume *more* than expected. Thus these solutions represents points at which the agent has maximum incentive to *conserve* electricity. Consequently, the point at which the peak incentive to waste is maximised is:

$$\omega_i^* = \frac{-1 + \mu_i \cdot \sqrt{\theta_i}}{\sqrt{\theta_i}} \tag{5.12}$$

which, when substituted into Equation 5.7, gives the condition for wasting as:

$$\frac{\theta_{max}^{3/4} \cdot \psi}{\sqrt{e} \cdot \pi^{1/4}} > f_r \tag{5.13}$$

However, to determine whether an agent has incentive to waste, it is still necessary to determine the upper bound on ψ. This is determined by a number of factors including the precision of the aggregator's priors. Moreover, the theoretical maximum savings is unbounded[1] and consequently, the *expected* ψ must be used:

$$\bar{\psi} = \frac{\lambda \cdot \bar{\Delta} \left(\boldsymbol{x}_{-i} \cup \{x_{a,i}\}, \boldsymbol{x}_a \right)}{S_{max} + \sum_{j \in N \setminus \{i\}} \bar{S}\left(x_j\right)} \tag{5.14}$$

where $\bar{\Delta}$ and $\bar{S}$ denote the expected savings function and the expected score respectively. Thus, assuming fixed market costs and lambdas, the peak expected ψ for a two agent scenario is:

$$\bar{\psi}^* = \frac{\lambda \cdot \bar{\Delta} \left(\mu_{a,i} + \mu_j, \left(\theta_{a,i}^{-1} + \theta_j^{*\,-1} \right)^{-1}, \boldsymbol{x}_a \right)}{S_{max} + \sum_{j \in N \setminus \{i\}} \bar{S} \left(\mu_j, \theta_j^* \right)} \tag{5.15}$$

where θ_j^* is the solution to:

$$\frac{\mathrm{d}}{\mathrm{d}\theta_j} \frac{\lambda \cdot \bar{\Delta} \left(\mu_{a,i} + \mu_j, \left(\theta_{a,i}^{-1} + \theta_j^{-1} \right)^{-1}, \boldsymbol{x}_a \right)}{S_{max} + \sum_{j \in N \setminus \{i\}} \bar{S} \left(\mu_j, \theta_j \right)} = 0 \tag{5.16}$$

The analytical solution to this problem quickly becomes intractable to compute. However, numerical analysis can be used in order to find the value of $\bar{\psi}^*$ given a set of parameters. Running this analysis on the parameters from Table 5.1 finds $\bar{\psi}^* = 1.52$.

[1]For example, consider the case in which $\chi\left(\boldsymbol{x}_a\right) \to \infty$ and $\chi\left(\boldsymbol{x}\right) = \omega$.

Thus the condition for wasting becomes:

$$\frac{\theta_{max}^{3/4} \cdot \bar{\psi}^*}{\sqrt{e} \cdot \pi^{1/4}} > f_r \tag{5.17}$$

with

$$\theta_{max} = 1500.00 \quad \bar{\psi}^* = 1.52 \quad f_r = 150.00$$

$$166.91 > 150.00 \tag{5.18}$$

Consequently, under the sum of others' plus max mechanism there *is* an incentive to waste when agents report the maximum precision and consume an amount of electricity exactly one standard deviation away from their reported mean. This can be fairly easily resolved through manipulation of the retail price, as is discussed next.

5.2.4 Removing the Incentive to Waste

In order to ensure that at no point is there any incentive to waste, it must be ensured that in the event of the worst case (that is, when the change in score for wasting is at its maximum and $\psi = \bar{\psi}^*$) Equation 5.17 is false. That is, the utility gain from the agent wasting electricity must be less than the cost of that electricity. Of course, a simple method of doing this, if $\bar{\psi}^*$ can be bounded, is to simply set f_r such that Equation 5.17 is always true. This is by far the simplest method, and retains budget balance and individual rationality for the aggregator. Further methods might be to modify the agents' reward functions such that the incentive to waste is lost. However, doing so is likely to be non-trivial while maintaining both budget balance and incentive compatibility.

However, using the simple method of real-time updating of f_r maintains both individual rationality and budget balance. This can be seen as the amount paid by agents ($f_r \cdot \omega_i$) is not dependent upon their reports. Therefore, misreporting does not affect their reward. Furthermore, it can be seen that budget balance is unaffected because the savings made by the aggregator are dependent upon the *market prices* (f, δ^b, and δ^s) and not the retail price, f_r.

One disadvantage of using this technique is that it can potentially result in large retail prices for the home agents. Consequently, it is not clear how agents will behave under such a model as determining this would require a new model of home agents which incorporates their utility for every unit of electricity consumed. With fixed retail pricing, this problem can be side-stepped by assuming that the agents know the retail price and still place their loads. That is, when placing loads ‘traditionally’ (i.e. under a fixed price or time of use tariff), loads are placed on the grid by the occupants knowing what their cost per unit of electricity is. Consequently, the UK data described within Section 4.4.1.1 can be used. However, with variable retail pricing, demand will also be dependent upon the time at which a load is placed (or more specifically, the cost of electricity at the time at which it is placed). This is not taken into account in the dataset from Section 4.4.1.1 and further requires the development of a model of utility per unit of electricity in order for agents to be able to strategies over what loads they place. Thus, only the behaviours of agents with fixed retail prices are taken into account in this thesis.

5.3 Conclusions from this Chapter

This chapter has discussed the theoretical and empirical analysis of the uniform and sum of others’ plus max mechanisms under scenarios in which the home agents are able to make themselves harder to predict. Crucially, by making themselves harder to predict, the agents decrease the precision of the aggregator’s priors. Consequently, the cost incurred by the aggregator when purchasing electricity using its own priors is increased.

It was shown that simply dividing the savings made by the aggregator is in fact *detrimental* to the aggregator as agents are encouraged to make themselves harder to predict in order to increase the perceived savings made by the aggregator and thus increase their rewards. Doing so rewards agents with a large utility but this comes at the cost of the aggregator, who makes a *loss* in utility compared to if it were to simply use no mechanism.

Moreover, it was shown that using the sum of others’ plus max mechanism encouraged agents to make themselves easier to predict. This, combined with the fact that agents are also incentivised to produce precise predictions, resulted in precise aggregate beliefs, which caused an *improvement* in utility to *both* the aggregator and the home agents. These results culminated in an increase in social welfare for the sum of others’ plus max mechanism; An increase which surpassed that of the uniform mechanism, which for small numbers of agents actually caused a loss of social welfare compared to using no mechanism.

In conclusion it was found that the sum of others’ plus max mechanism provided far more beneficial incentives to the home agents than the uniform mechanism did. Furthermore, it was shown that the uniform mechanism would not be implemented by a rational profit-maximising aggregator as doing so causes the aggregator to make a loss in utility.

This concludes the main body of work for this thesis. The next chapter draws the thesis to a close by summarising the work done and summarising the conclusions from said work. In addition, the next chapter discusses possible avenues of future research.

Chapter 6

Conclusions and Future Work

This thesis has discussed the problem of aggregate demand prediction within the smart grid. This chapter will draw the thesis to an end by discussing the conclusions from this work as well as possible directions of future work that can incorporate the work from this thesis.

6.1 Conclusions

This thesis has discussed the aggregate demand prediction problem within the smart grid. In this problem, an aggregator agent is tasked with purchasing electricity for a set of consumers. These consumers are assumed to be homes, and each home is represented by its own home agent. In order to purchase electricity, the aggregator must make use of one of two markets. In order to purchase electricity ahead of time, the aggregator must purchase electricity from the forward markets. However, if the aggregator does not purchase the amount actually consumed by the home agents in the forward market, it must buy or sell any imbalance in real time in the balancing markets. Prices are set to encourage accurate prediction of loads, and as such, trading electricity in the balancing market is more costly than doing so in the forward markets. Consequently, in order to minimise the cost of electricity, the aggregator must make a prediction as to how much electricity the homes will consume. To do this, the aggregator has some prior knowledge about the consumption of each home, which it gains through statistical analysis of that home's historical loads. However, much more precise information is available within each home regarding how that home will consume in the future. For example, information regarding devices within the home, the home's occupants' calendars and comfort levels etc. can be collected in order to give a precise estimate of how that home will consume in the future. The home agents can gather this information in order to send to the aggregator such that the aggregator is able to make a more optimal purchase of electricity from the forward markets. In doing this, the aggregator makes some amount of savings.

However, the home agents incur a cost by gathering this information. Consequently, the aggregator must make a payment to the home agents in order to encourage them to gather and report their information to the aggregator. In this work, it is proposed that the aggregator distributes a fraction of the *savings* made by the aggregator to the home agents. However, this is not trivial. The payments to the agents must be individually rational such that the agents have an incentive to participate; incentive compatible, such that the agents report their information truthfully to the aggregator; and budget balanced such that the aggregator does not make a loss through the use of the payment mechanism (in which case, the aggregator may as well simply use its priors). It is on the problem of designing such a payment mechanism which this thesis focuses.

In more detail, it is assumed within this thesis that the home agents may try to game the mechanism in a number of ways. For one, agents may misreport their information. For example, they may claim that their reports contain more information than they do in reality. In addition, agents may make themselves harder to predict. That is, they may deliberately choose to make the loads they place upon the grid more volatile such that the aggregator's prior information (based on historical records) is less precise. Finally, agents may choose to waste electricity such that their real consumption maximises their expected reward. For example, an agent who knows it will be paid a fixed amount only if it consumes the expected value of the distribution it reported will likely waste electricity if it finds that it will consume less than what it had originally reported.

As part of the solution, the use of scoring rules is proposed. These are functions which map probabilistic reports and events to a real-valued score. Of particular interest are *proper* and *strictly proper* scoring rules, which are maximised in expectation (uniquely in the latter case) when an agents report their information truthfully. These rules have been developed such that they do not incentivise 'undesirable' behaviour, (Shi et al., 2009; Bacon et al., 2012). However, they have not taken into account either budget balance of payments to agents nor the fact that agents may make themselves harder to predict. Indeed, budget balance is something which is lacking from prior literature on scoring rules within the continuous domain. Usually, it is simply assumed that agents are somehow incentivised to maximise their score, in which case payments are independent of their actual scores, or budget balance is simply not taken into account whatsoever.

Furthermore, prior literature in mechanism design in general was found to be lacking in that it did not take into account incentives to manipulate by agents making themselves harder to predict or fixing their consumptions. Moreover, prior work in mechanism design has used fixed budgets whereas the budget here is dependent upon the agents own actions. Further to that, it is often assumed that a number of agents are available from which the aggregator can request information, and that the aggregator has no prior information. The work in this thesis assumes that while there are a number of agents from whom the aggregator can request information, each agent predicts an individual event and it is the *cumulative* effect of these events which the aggregator must plan for.

Consequently, prior works in mechanism design which reward agents by comparing them to one-another cannot be used as the agents are predicting different events.

To this end, the work in this thesis develops a scoring rule named *sum of others' plus max*, which can be used to distribute payments in an individually rational, dominant strategy incentive compatible and weakly budget balanced manner. Furthermore, when combined with the use of a budget based upon savings, the *sum of others' plus max mechanism* can be used to distribute payments to agents such that their reward is directly affected by their reports. Moreover, under such a mechanism, the rewards are based upon the *benefit* the agents provided to the aggregator. This mechanism was compared to the *uniform* mechanism, an individually rational, strictly budget balanced but only *Nash* incentive compatible mechanism which distributes the allocated savings evenly among the home agents. In addition it was compared to using no mechanism at all (i.e. the aggregator simply uses its priors and requests no information from the home agents).

Results showed that the scoring rule-based mechanism provides better incentives for the home agents to produce precise reports, which in turn increases the expected utility of the aggregator and reduces the aggregator's risk in comparison to the uniform mechanism and in comparison to using no mechanism. Furthermore, results showed that the uniform mechanism encourages agents to make their loads far more volatile such that they were harder to predict, whereas the sum of others' plus max mechanism incentivises agents to *reduce* the volatility of their loads. The effect of this is that under the uniform mechanism, the aggregator makes a large loss in utility in comparison to when using no mechanism, resulting from the reduced precision of the aggregator's priors. Under the same conditions, the sum of others' plus max mechanism causes an increase in expected utility for the aggregator and a decrease in risk as it incentivises agents to make themselves even easier to predict by reducing the volatility of their loads. As well as reducing the cost of electricity, this behaviour is beneficial to the grid as the amount of spinning reserve generation can be reduced, and generation can be more optimally planned.

However, in terms of wasting electricity, it was found that the sum of others' plus max mechanism may encourage home agents to waste electricity if they find that they are going to consume less than they predicted, whereas this is not the case for the uniform mechanism. Conditions for this wasting behaviour were discussed, and it was shown that this behaviour can be easily mitigated through manipulation of the retail price (the price at which the home agents purchase electricity from the aggregator). The possibility of using variable retail prices was discussed. However, implementing a simulation of such a scenario would require a large amount of complexity to be added into the model from this thesis, including a model of the homes' value of each unit of electricity consumed (such that they are able to determine whether it is worth placing extra demand). Consequently, this was left for future work, a discussion of which appears in the next section.

6.2 Future Work

While the work in this thesis has provided extensive analysis to the design of mechanisms for the rewarding of agents supplying predictive information within the smart grid, there are still some steps that are required before such mechanisms can be deployed to real homes.

First, home users are unlikely to be well versed in the theory of mechanism design and game theory. Therefore, when interacting with the agents, a lot of the complexity must be hidden away. This is already an active field of research within the human-agent interaction community. In particular, work from Seuken et al. (2012) addresses the development of user interfaces that hide the complexity of markets. Furthermore, work (for example in Rogers et al. (2011)) is being carried out in the development of smart devices within homes that allows users to easily input their preferences to allow agents to control devices within the property. The future work required to apply the mechanisms from this thesis into a real world scenario lies at the intersection of these two areas; not only must an interface be designed such that users gain an understanding of their actions upon their rewards, but also it must somehow learn the users' preferences such that the home agent is able to manipulate the loads placed by that house without affecting the occupants. Given this work, user trials can be carried out in order to test how the use of the mechanisms in this thesis along with the appropriate interface can change the behaviour of humans, who are known not to behave rationally. It would be particularly interesting to see the results of a user-based experiment into the effect of the uniform and sum of others' plus max mechanisms on the users' load variance manipulation. For example, Chapter 5 showed that under the assumption of rational agents the uniform mechanism caused users' loads to become very unpredictable. It would be interesting to see whether or not this holds for humans.

Another direction for future work is to relax the assumption that agents' consumption information is private to themselves. By relaxing this assumption it is possible to have the home agents make predictions about their neighbourhood's future consumption. Such a mechanism could make use of prediction markets wherein agents bid on the future consumption of their neighbourhood in an effort to reduce the cost of their own electricity. However, it is not clear whether the aggregator, or even the social welfare will benefit from such a scenario.

A limitation of the work in this thesis which negatively affects the utility to the aggregator is that the aggregator purchases information from all agents, regardless of the informational value of their reports. That is, even if the agent reports information of the same precision as the aggregator's priors, that agent will still receive a reward from the mechanism (albeit a smaller reward than it would have received had it reported information of higher precision). An extension of this work could look into scenarios in which the aggregator does not purchase information from all agents, but instead only purchases information from the agents whose reports are likely to make the most significant improvement to the aggregate belief. This is similar to the work in Papakonstantinou et al. (2011) in which information is gathered from agents with unknown costs. However, in their work it is assumed that there is sufficient time and resources for multiple stages of a mechanism (in which costs and precisions are elicited, the agents are selected and then the information is elicited and fused from those agents). Consider a scenario such as making a prediction before gate closure in which only a single ask-and-gather stage is feasible. How should the aggregator choose the agents from which it will purchase information given the information already held by the aggregator? This problem is not only present within the aggregation scenario of this thesis, but also in other fields. For example, from the field of sensor networks, a base station might be making predictions of future conditions given information held by itself and a set of sensor nodes. This time, the limitation to a single-stage game is not due necessarily to time constraints, but due to issues of power conservation. Each time a node sends and receives data, that node consumes energy of which it only has a finite amount. Thus, the base station must request information only from those agents whom it believes will maximise the change to the base station's prediction.

Furthermore, if the mechanism could be extended to ensure that the agents don't fuse into their reports information that is already held by the aggregator, the aggregator could fuse information from the home agents with its own. Thus information from the home agents could always be used to improve the aggregator's beliefs, regardless of the home agents' precision. Moreover, this would reduce the computational burden on the home agents, albeit at the cost of the aggregator, to which this computational load will be shifted.

From a mechanism design perspective, it would be interesting to determine whether or not a strictly budget balanced scoring rule-based mechanism could be found. It is not clear whether this is in fact possible. It would be interesting to prove theoretically whether or not there exists a mechanism for the model described in Section 3.2 which allocates payments that are budget balanced, dominant strategy incentive compatible and strictly budget balanced. It is very likely that such a mechanism does not exist, as enforcing strict budget balance will likely break dominant strategy incentive compatibility as it did in the percentage contribution mechanism discussed in Section 4.2. Another potential extension along these lines could be for the aggregator to send information to the agents for them to fuse into their own reports.

Of course, a large amount of effort is currently being invested into the invention of efficient and cost effective storage solutions for electricity. This eventuality is not modelled by the work in this thesis. If storage were available, the home agents' utility functions would differ significantly to those that are used within this thesis. For one, in terms of agents wasting electricity in order to ensure they consume as predicted, agents might still gain positive utility for consuming their extra units of electricity and placing them into storage. This is because rather than simply wasting the electricity, the agents will gain some amount of utility based on their expected cost of purchasing and storing the electricity now compared to purchasing and using the electricity later. In such a situation it may even be advantageous to the grid as a whole for agents to behave in this way (i.e. to store electricity they predicted they would consume for later use). It would be of great benefit to the future of research into mechanism design for the smart grid to develop mechanisms that took into account storage.

Appendix A

Derivation of the Logarithmic, Quadratic and Spherical Scoring Rules

This appendix discusses the derivation of the logarithmic and quadratic and spherical scoring rules. This uses the characterisations by McCarthy (1956), Hendrickson and Buehler (1971), Savage (1971) and Schervish (1989) as is discussed in Sections 2.6.1.1 and 2.6.2.

A discrete scoring rule, $S : \mathbb{P} \times \Omega \to \mathbb{R}$ is proper iff:

$$S(\boldsymbol{p}, \omega) = G(\boldsymbol{p}) - \big\langle G'(\boldsymbol{p}), \boldsymbol{p} \big\rangle + G'_i(\boldsymbol{p}) \ \forall i \in \{1, \cdots, n\} \tag{A.1}$$

where the term $\langle G'(\boldsymbol{p}), \boldsymbol{p} \rangle$ represents the dot product of the sub-derivative of G with respect to $\boldsymbol{p}$, $G'(\boldsymbol{p})$, and $\boldsymbol{p}$, and $G'_i(\boldsymbol{p})$ is the value of $G'(\boldsymbol{p})$ at event i.

Furthermore, a continuous scoring rule, $S : \mathbb{P} \times \Omega \to \mathbb{R}$, is proper iff S is of the form:

$$S(p, \omega) = G(p) - \int G^*(p, \omega)\, \mathrm{d}P(\omega) + G^*(p, \omega)$$

where, $P(\omega)$ is the cumulative distribution function at ω. Given that the probability density function, $p(\omega) = \frac{\mathrm{d}P(\omega)}{\mathrm{d}\omega}$ and thus $\mathrm{d}P(\omega) = p(\omega)\ \mathrm{d}\omega$, this becomes:

$$S(p, \omega) = G(p) - \int p(\omega)\, G^*(p, \omega)\ \mathrm{d}\omega + G^*(p, \omega)$$

It is shown in Theorem 2.1 that $\mathbb{P}$ is convex, and therefore, for G to be convex, it is sufficient to show that (R. Tyrrell Rockafellar, 1970):

$$G(\theta x + (1 - \theta) y) \leq \theta G(x) + (1 - \theta) G(y),\ \forall\, x, y \in \mathbb{P},\ \forall\, \theta \in [0, 1]$$

The derivations of the continuous and discrete versions follow a similar technique. Consequently, only the derivations for the discrete scoring rules are shown. This appendix begins by presenting the derivation of the logarithmic scoring rule, followed by the quadratic scoring rule and then finally the spherical scoring rule.

A.1 Derivation of the Logarithmic Scoring Rule

A.1.1 Expected Score Function

Begin with the expected score function

$$G(\boldsymbol{p}) = \boldsymbol{p} \cdot \log(\boldsymbol{p}) \tag{A.2}$$

Take the first derivative with respect to $\boldsymbol{p}$

$$G^*(\boldsymbol{p}) = 1 + \log(\boldsymbol{p}) \tag{A.3}$$

A.1.2 Deriving the Rule

Substituting into Equation A.1:

$$S(\boldsymbol{p}, \omega) = \boldsymbol{p} \cdot \log(\boldsymbol{p}) - \boldsymbol{p} \cdot (1 + \log(\boldsymbol{p})) + \mathbf{0}_{\boldsymbol{p}}(\omega) \cdot (1 + \log(\boldsymbol{p})) \tag{A.4}$$

where $\mathbf{0}_{\boldsymbol{p}}(\omega)$ is a vector of length $|\boldsymbol{p}|$ in which the ω^{th} element is 1 and all other elements are 0. Thus:

$$S(\boldsymbol{p}, \omega) = \boldsymbol{p} \cdot \log(\boldsymbol{p}) - \boldsymbol{p} \cdot \mathbf{1}_{\boldsymbol{p}} - \boldsymbol{p} \cdot \log(\boldsymbol{p}) + 1 + log(p_\omega) \tag{A.5}$$

where $\mathbf{1}_{\boldsymbol{p}}$ is a vector of length $|\boldsymbol{p}|$ comprised solely of 1s. $\boldsymbol{p} \cdot \mathbf{1}_{\boldsymbol{p}} = 1$ consequently, the terms cancel to give:

$$S(\boldsymbol{p}, \omega) = log(p_\omega). \tag{A.6}$$

A.2 Derivation of the Quadratic Scoring Rule

A.2.1 Expected Score Function

Begin with the expected score function

$$G(\boldsymbol{p}) = \boldsymbol{p} \cdot \boldsymbol{p} - 1 \tag{A.7}$$

The first derivative of which with respect to $\boldsymbol{p}$ is:

$$G^{*}(\boldsymbol{p}) = 2\boldsymbol{p}. \tag{A.8}$$

A.2.2 Deriving the Rule

Thus, substituting into Equation A.1 gives:

$$S(\boldsymbol{p}, \omega) = \boldsymbol{p} \cdot \boldsymbol{p} - 1 - 2\boldsymbol{p} \cdot \boldsymbol{p} + \mathbf{0}_{\boldsymbol{p}}(\omega) \cdot 2\boldsymbol{p} \tag{A.9}$$

Which simplifies to give the quadratic scoring rule minus one. Adding one maintains incentive compatibility, and results in the quadratic scoring rule.

$$S(\boldsymbol{p}, \omega) = 2p_{\omega} - \boldsymbol{p} \cdot \boldsymbol{p} \tag{A.10}$$

A.3 Derivation of the Spherical Scoring Rule

A.3.1 Expected Score Function

Begin with the expected score function:

$$G(\mathbf{p}) = \left(\sum_{j=1}^{n} p_j^2 \right)^{\frac{1}{2}} \tag{A.11}$$

and therefore

$$G'(\mathbf{p}) = \frac{\sum_{j=1}^{n} p_j}{\left(\sum_{j=1}^{m} p_j^2 \right)^{\frac{1}{2}}} \tag{A.12}$$

A.3.2 Deriving the Rule

Then, substituting these functions into Equation A.1 gives:

$$\begin{aligned}
S(\mathbf{p},\omega) &= \left(\sum_{i=1}^{n} p_i^2\right)^{\frac{1}{2}} - \mathbf{p} \cdot \left(\sum_{i=1}^{n} p_i\right) \cdot \left(\sum_{i=1}^{n} p_i^2\right)^{-\frac{1}{2}} + p_\omega \cdot \left(\sum_{i=1}^{n} p_i^2\right)^{-\frac{1}{2}} \\
&= \left(\sum_{i=1}^{n} p_i^2\right)^{\frac{1}{2}} - \left(\sum_{i=1}^{n} p_i^2\right) \cdot \left(\sum_{i=1}^{n} p_i^2\right)^{-\frac{1}{2}} + p_\omega \cdot \left(\sum_{i=1}^{n} p_i^2\right)^{-\frac{1}{2}} \\
&= \left(\sum_{i=1}^{n} p_i^2\right)^{\frac{1}{2}} - \left(\sum_{i=1}^{n} p_i^2\right)^{\frac{1}{2}} + p_\omega \cdot \left(\sum_{i=1}^{n} p_i^2\right)^{-\frac{1}{2}} \\
&= p_\omega \cdot \left(\sum_{i=1}^{n} p_i^2\right)^{-\frac{1}{2}}
\end{aligned} \tag{A.13}$$

Appendix B

Approximating the Sum of Log-Normal Distributions using Gaussians

The exact sum of numerous log-normal distributions is computationally hard to compute. Doing so requires the computation of the convolution of the probability density functions of the individual distributions that are to be summed. However, for carefully constrained log-normal distributions, it is possible to sum the distributions as Gaussian distributions. This appendix provides a discussion on this approximation method.

B.1 Hellinger Distance

There are a number of ways of measuring the difference between distributions. The method within used within this appendix is the Hellinger distance of two distributions P and Q, defined as:

$$H^2(P, Q) = 1 - \int_{-\infty}^{\infty} \sqrt{P(x) \cdot Q(x)}\, \mathrm{d}x. \tag{B.1}$$

The range of the Hellinger distance is between 0 and 1, with values approaching 1 indicating a larger difference between the two distributions. More specifically, a Hellinger distance of 1 indicates that all outcomes indicated as having a probability of 1 by one distribution have a probability of 0 according to the other distribution. Moreover, a Hellinger distance of 0 indicates that the two distributions are equivalent.

Hellinger distance was used as a metric within this appendix due to the fact that it is symmetric. That is, the metric is concerned only with the *distance* between two

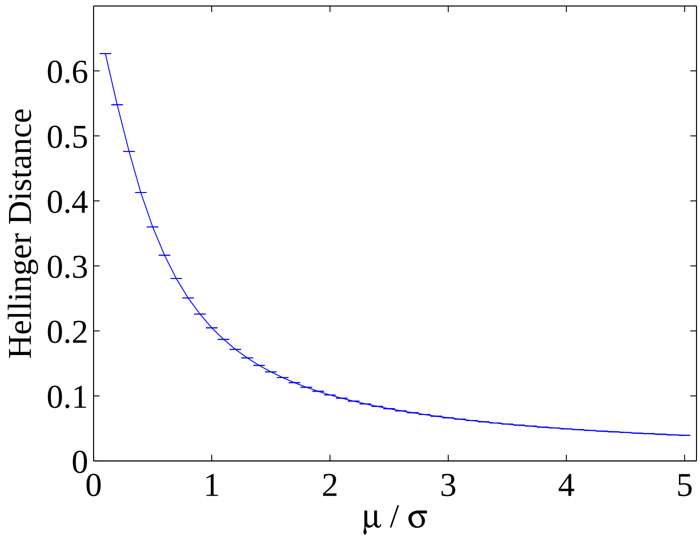

FIGURE B.1: The average Hellinger distance from 10 samples of the sum of 5 log-normal distributions with means that is are x times greater than their standard deviation, $\sigma = 1$.

distributions. Other measures, such as the commonly used KL divergence, $D_{KL}(P, Q)$, do not exhibit this property as they measure the information lost or gained in using one distribution to approximate another. Consequently $D_{KL}(P, Q) \neq D_{KL}(Q, P)$. The analysis here is concerned only with how well the probability density functions of the resulting sums are matched, and therefore, the Hellinger distance is used.

B.2 Approximation of the Sum of Log-Normal Distributions using Gaussians

How well the sum of a number of log-normal distributions is approximated as a Gaussian is highly dependent upon the parameters of the log normal distributions. Figure B.1 shows the result from which 5 log normal distributions with random means and a fixed standard deviation of $\sigma = 1$ are summed. The distributions' means are all equal to the value of the x axis multiplied by $\sigma = 1$. This therefore shows how well the Gaussian approximation of the sum of log-normal distributions with the same means and standard deviations matches the real convolution of said distributions. It can be seen that as the means become further from 0 the approximation improves. This is due to the fact that log-normal distributions have a lower bound of 0 whereas Gaussians are unbounded. Consequently, the redistribution of probability densities such that events less than zero are impossible disfigures the log-normal distribution compared to the Gaussian.

Further to that, Figure B.2 shows the results from an experiment similar to the one above except where the means of the distributions are sampled from a log-normal distribution

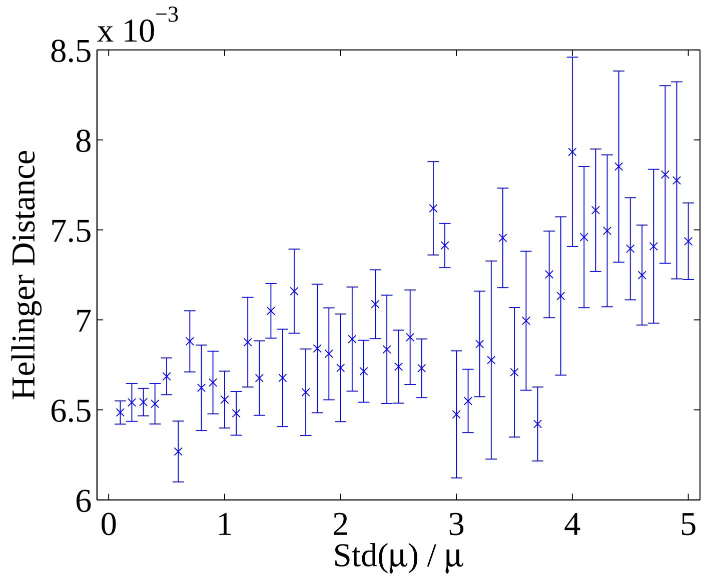

FIGURE B.2: The average Hellinger distance from 10 samples of the sum of 5 log-normal distributions with means randomly distributed around a point far from zero (30) where the standard deviation of the means of the distributions is, $x \cdot 30$.

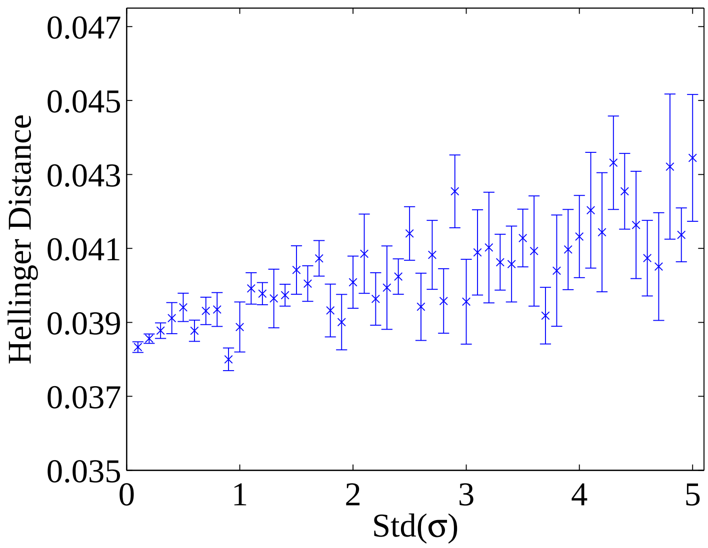

FIGURE B.3: The average Hellinger distance from 10 samples of the sum of 5 log-normal distributions with mean 50 and standard deviations randomly distributed around a mean 10 with standard deviation, $Std(\sigma)$.

with an expected value of 30 and a standard deviation of the x value times 30. That is, as x increases, the difference between the *means* of the log-normal distributions to be summed increases. It can be seen that the distance between the approximation and the exact solution is of the order 10^{-3}. Consequently, the approximation appears to work well even for the sum of distributions with diverse means, provided the means are sufficiently far from zero.

Finally, Figure B.3 shows the results from an experiment similar to that above, but in which all distributions have the same mean of 50, chosen to be far from zero compared to their standard deviations, and the standard deviations of the distributions are sampled from a log-normal distribution with mean 10 and standard deviation x. Again, this shows that the error is very small, even with distributions with diverse standard deviations.

B.3 Conclusions

Given the results shown in the previous section, the Gaussian approximation of the sum of log-normal distributions appears to give relatively accurate results provided the means of the distributions to be summed are far enough away from zero.

Appendix C

Home Consumption Data

This appendix contains the histograms of the consumptions of the 16 houses whose data were used in the simulations in the simulations of Chapters 4 and 5.

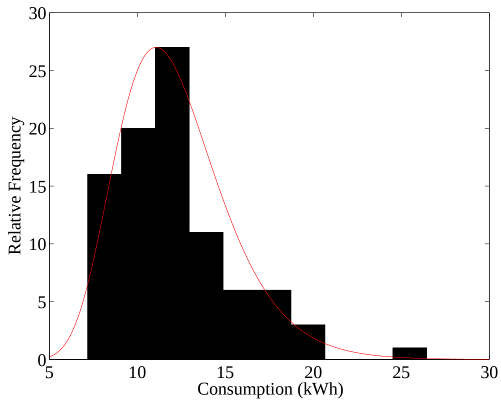

(a) House 1

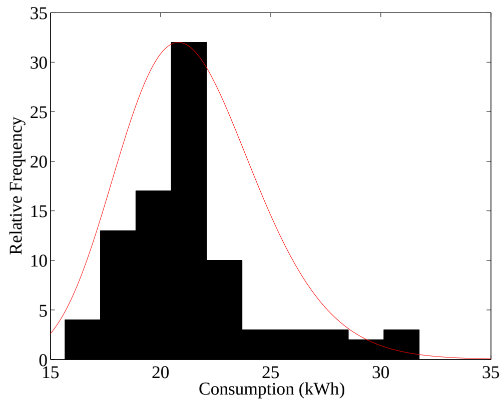

(b) House 2

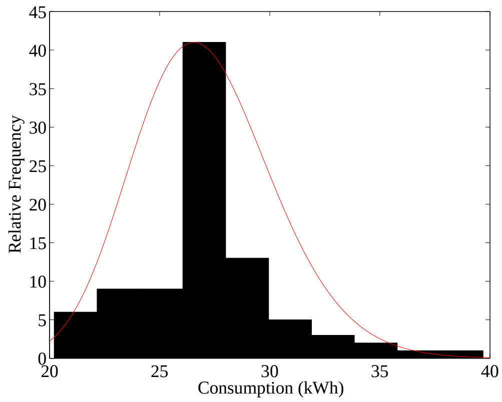

(c) House 3

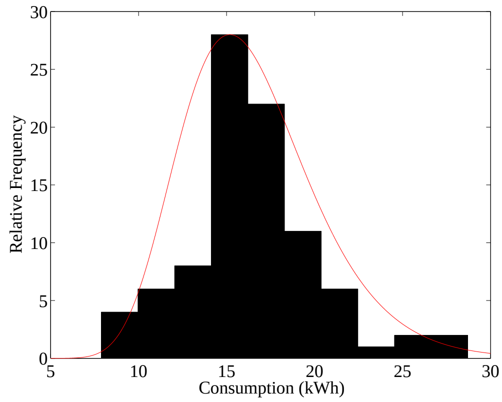

(d) House 4

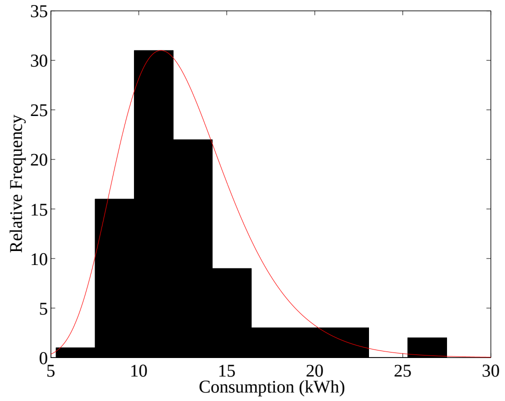

(e) House 5

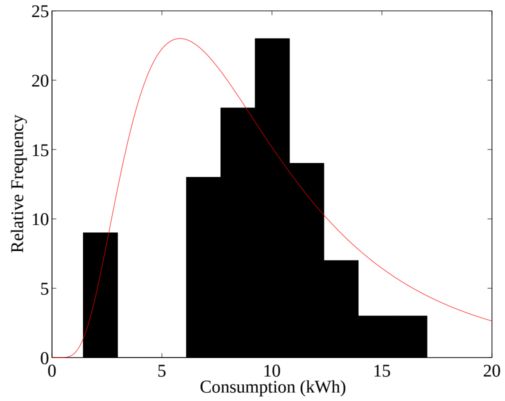

(f) House 6

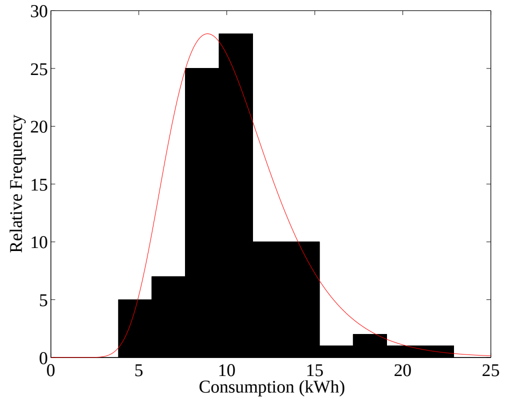

(g) House 7

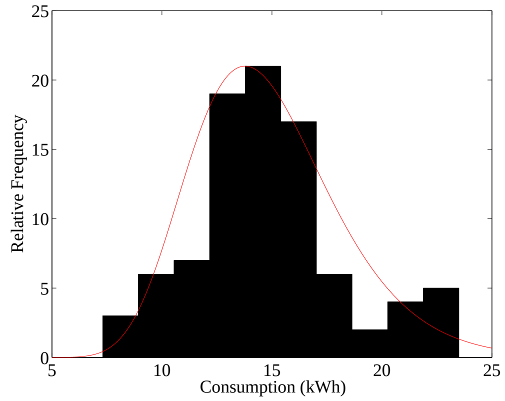

(h) House 8

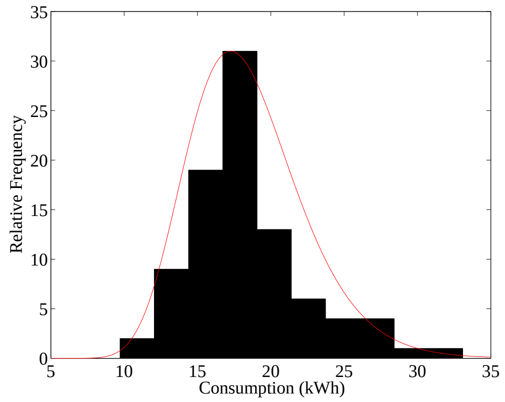

(i) House 9

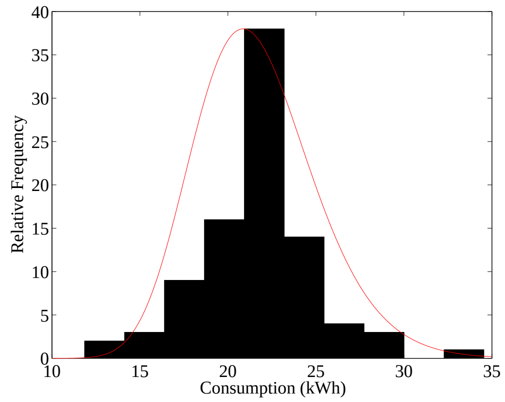

(j) House 10

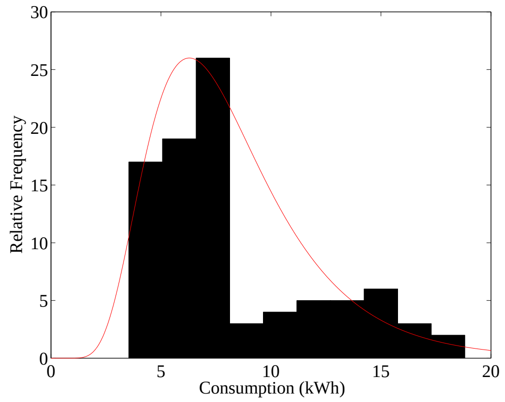

(k) House 11

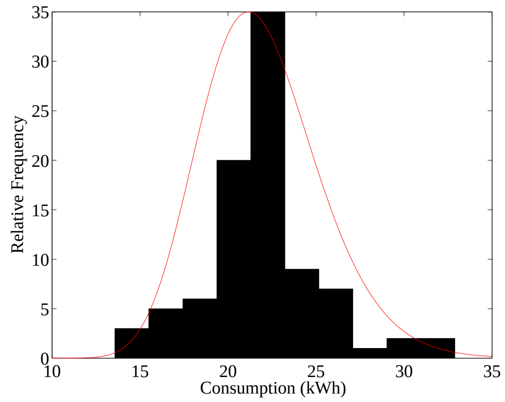

(l) House 12

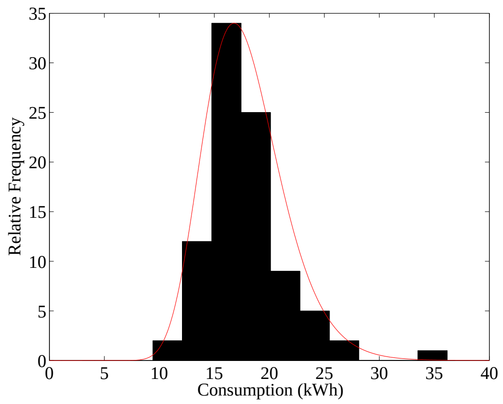

(m) House 13

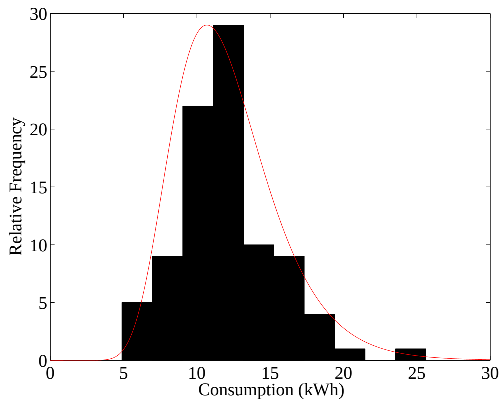

(n) House 14

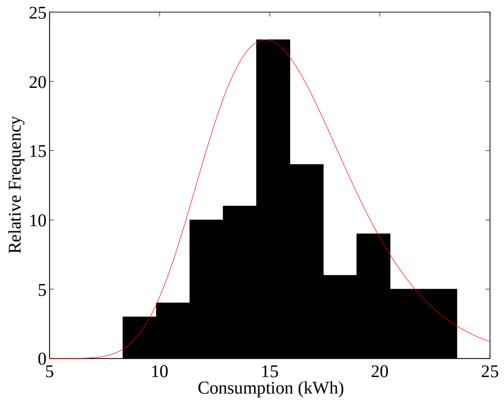

(o) House 15

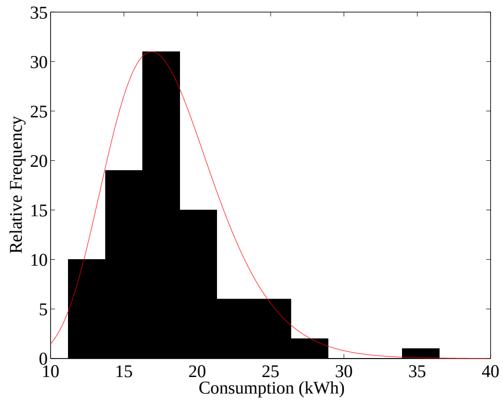

(p) House 16

Appendix D

The Sampled Peak Finding Algorithm

This appendix describes the peak finding algorithm that was used within the simulations in this thesis. This algorithm was used as opposed to other algorithms such as hill-climbing algorithms as it is guaranteed to converge upon a peak, whereas hill-climbing algorithms might constantly overshoot the peak. A graphical representation of the algorithm described in Listing D.1 is shown in Figure D.1.

D.1 JAVA Code of the Algorithm

```
public static double[] maximise(TwoParameterFunction function, double xMin,
                        double xMax, double yMin, double yMax, double tolerance)
{
        // The distance between the upper and lower bounds of the search
        // space is at least as small as the threshold. Return the centre
        // of the search space.
        if(xMax - xMin <= tolerance && yMax - yMin <= tolerance)
        {
                return new double[]{ (xMax + xMin)/2.0, (yMax + yMin)/2.0};
        }

        int samples = 3;

        double bestx = Double.NaN;
        double besty = Double.NaN;
        double bestv = Double.MIN_VALUE;

        double dx = (xMax - xMin) / samples;
        double dy = (yMax - yMin) / samples;

        double v = 0;
        for(double x = xMin; x <= xMax; x += dx)
        {
                for(double y = yMin; y <= yMax; y += dy)
```

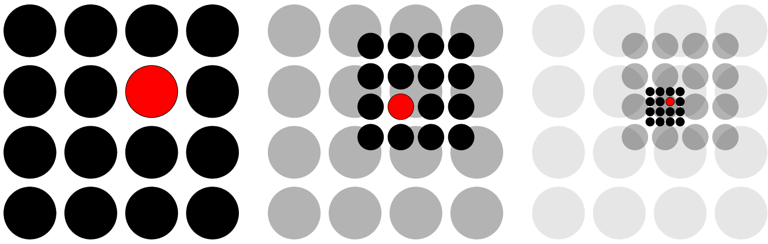

FIGURE D.1: A graphical representation of how the peak is found using the algorithm in Listing D.1. Black points represent points being checked in the current iteration, grey points represent the points that have already been checked and red points represent the highest point found in that particular iteration.

```
            {
                    v = function.getValueAt(x, y);

                    if( v > bestv )
                    {
                            bestx = x;
                            besty = y;
                            bestv = v;
                    }
            }
    }

    // Look around the best point
    return maximise(function,bestx - dx, bestx + dx, besty - dy,
                    besty + dy, tolerance);
}
```

LISTING D.1: Optimisation algorithm used to find the agents' optimal strategies.

D.2 Complexity

After every iteration of the search, the size of the search space decreases by a factor of

$$\frac{2}{samples}$$

Therefore, given a starting space of width $startWidth$, after iteration i, the search space has the width:

$$width_i = startWidth \cdot \left(\frac{2}{samples}\right)^i = startWidth \cdot \left(\frac{samples}{2}\right)^{-i}$$

The algorithm continues until the $width_i = thresholdWidth$. Therefore, the algorithm will stop when

$$thresholdWidth = startWidth \cdot \left(\frac{samples}{2}\right)^{-i}$$

That is, after iteration:

$$i = -\log_{\frac{samples}{2}}\left(\frac{thresholdWidth}{startWidth}\right)$$

Therefore the complexity of the search algorithm with respect to the search space is:

$$O\left(\log_{\frac{samples}{2}}\left(\frac{startWidth}{thresholdWidth}\right)\right)$$

Moreover, within each iteration, the complexity is dependent upon the number of samples performed. The complexity of each iteration is thus $O\left(samples^2\right)$.

Therefore the overall complexity of the algorithm is:

$$O\left(samples^2 \cdot \log_{\frac{samples}{2}}\left(\frac{startWidth}{thresholdWidth}\right)\right)$$

References

David F. Bacon, Yiling Chen, Ian Kash, David C. Parkes, Malvika Rao, and Manu Sridharan. Predicting Your Own Effort. *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012), Valencia, Spain*, 2012.

J Eric Bickel. Some Comparisons among Quadratic, Spherical and Logarithmic Scoring Rules. *Decision Analysis*, 4(2):49–65, 2007.

Severin Borenstein. The Trouble with Electricity Markets: Understanding California's Restructuring Disaster. *Journal of Economic Perspectives*, 16(1):191–211, 2002.

Craig Boutilier. Eliciting Forecasts from Self-interested Experts: Scoring Rules for Decision Makers. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, pages 737—784, Valencia, Spain, 2012.

Glenn Brier. Verification of Forecasts in Terms of Probability. *Monthly Weather Review*, 78(1):1–3, 1950.

E. Carpaneto and G. Chicco. Probabilistic Characterisation of the Aggregated Residential Load Patterns. *IET Generation, Transmission & Distribution*, 2(3):373, 2008.

Cristiano Castelfranchi. Guarantees for Autonomy in Cognitive Agent Architecture. In Michael Wooldridge and Nicholas R Jennings, editors, *Intelligent Agents*, Lecture Notes in Computer Science, pages 56–70. Springer Berlin / Heidelberg, 1995.

Georgios Chalkiadakis, Valentin Robu, Ramachandra Kota, Alex Rogers, and Nicholas R. Jennings. Cooperatives of distributed energy resources for efficient virtual power plants. In *The 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2011)*, pages 787–794, May 2011.

Vincent Conitzer and Tuomas Sandholm. New complexity results about Nash equilibria. *Games and Economic Behavior*, 63(2):621—641, July 2008.

Sarah Darby. Making it Obvious: Designing Feedback into Energy Consumption. In *2nd International Conference on Energy Efficiency in Household Appliances and Lighting*. Italian Association of Energy Economists / EC-SAVE programme, 2001.

Sarah Darby. The Effeciveness of Feedback on Energy Consumption. Technical report, Environmental Change Institute – University of Oxford, 2006.

Partha S Dasgupta, Peter J Hammond, and Eric S Maskin. The Implementation of Social Choice Rules: Some General Results on Incentive Compatibility. *Review of Economic Studies*, 46(2):185 – 216, 1979.

Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The complexity of computing a Nash equilibrium. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of Computing - STOC '06*, pages 71—78, Seattle, Washington, USA, May 2006.

Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The complexity of computing a Nash equilibrium. *Communications of the ACM*, 52(2):89, February 2009.

Department of Energy and Climate Change. *Digest of United Kingdom Energy Statistics 2009*. Stationery Office Books (TSO), 2009a.

Department of Energy and Climate Change. *The UK Low Carbon Transition Plan: National Strategy for Climate and Energy*. The Stationary Office, 2009b.

Roger Dettmer. Living with NETA [New Electricity Trading Arrangements]. *IEE Review*, 48(4):32–36, 2002.

Electricity Networks Strategy Group. Smarter Grids: The Opportunity. Technical report, Department of Energy & Climate Change UK, 2009.

European Commission and Directorate-General for Mobility and Transport. Interconnecting Europe: New Perspectives for Trans-European Energy Networks. Technical report, {Directorate-General for Mobility and Transport}, 2008.

Daniel Friedman. Effective Scoring Rules for Probabilistic Forecasts. *Management Science*, 29(4):447–454, 1983.

Drew Fudenberg. *Game Theory*. McGraw-Hill, New York, 1991.

Drew Fudenberg and David K Levine. *The Theory of Learning in Games (Economic Learning and Social Evolution)*. The MIT Press, 1998. ISBN 0262061945.

Robert Galvin and Kurt Yeager. *PERFECT POWER : How the Microgrid Revolution Will Unleash Cleaner, Greener, More Abundant Energy*. McGraw-Hill, 2008.

Enrico Gerding, Kate Larson, and Nicholas Jennings. Eliciting Expert Advice in Service-Oriented Computing. In *Proceedings of the 11th International Workshop on Agent-Mediated Electronic Commerce (AMEC 2009)*, pages 29–42, 2009.

Tilmann Gneiting and Adrian E Raftery. Strictly Proper Scoring Rules, Prediction, and Estimation. Technical report, Department of Statistics, University of Washington, 2007.

Irving Good. Rational Decisions. *Journal of the Royal Statistical Society. Series B (Methodological)*, 14(1):107–114, 1952.

John C. Harsanyi. Games with incomplete information played by "bayesian" players, i-iii. part i. the basic model. *Management Science*, 14(3):pp. 159–182, 1967.

John C. Harsanyi. Games with incomplete information played by "bayesian" players, i-iii. part ii. bayesian equilibrium points. *Management Science*, 14(5):pp. 320–334, 1968.

Arlo Hendrickson and Robert Buehler. Proper Scores for Probability Forecasters. *The Annals of Mathematical Statistics*, 42(6):1916—1921, 1971.

R. Herman and J.J. Kritzinger. The statistical description of grouped domestic electrical load currents. *Electric Power Systems Research*, 27(1):43–48, May 1993. ISSN 03787796.

Serena Hesmondhalgh. Is NETA the Blueprint for Wholesale Electricity Trading Arrangements of the Future? *IEEE Transactions on Power Systems*, 18:548–554, 2003.

S.W. Heunis and R. Herman. A probabilistic model for residential consumer loads. *IEEE Transactions on Power Systems*, 17(3):621–625, August 2002. ISSN 0885-8950.

Frederick S. Hillier and Gerald J. Lieberman. *Introduction to Operations Research*. McGraw-Hill, 8^{th} edition, 2005.

John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behaviour*. Princeton University Press, 1944.

Andreu Mas-Colell, Michael D Whinston, and Jerry R Green. *Microeconomic Theory*. McGraw-Hill, New York, 1995.

James Matheson and Robert Winkler. Scoring Rules for Continuous Probability Distributions. *Management Science*, 22(10):1087–1096, 1976.

John McCarthy. Measures of The Value of Information. *Proceedings of the National Academy of Sciences of the United States of America*, 42(9):654–655, 1956.

D.H.O. McQueen, P.R. Hyland, and S.J. Watson. Monte Carlo Simulation of Residential Electricity Demand for Forecasting Maximum Demand on Distribution Networks. *IEEE Transactions on Power Systems*, 19(3):1685—1689, August 2004.

Nolan Miller, Paul Resnick, and Richard Zeckhauser. Eliciting Informative Feedback: The Peer-Prediction Method. *Management Science*, 51(9):1359–1373, 2005.

Allan Murphy. Hedging and Skill Scores for Probability Forecasts. *Journal of Applied Meteorology*, 12(1):215–223, 1973.

Roger B Myerson. Incentive Compatibility and the Bargaining Problem. *Econometrica*, 47(1):61 – 73, 1979.

Roger B. Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, 1991.

Victor Naroditskiy and Amy Greenwald. Using Iterated Best-Response to Find Bayes-Nash Equilibria in Auctions. In *Proceedings of the 22nd national conference on Artificial intelligence - Volume 2*, pages 1894—1895. AAAI Press, 2007.

John Nash. Equilibrium Points in n-Person Games. *Proceedings of the National Academy of Sciences of the United States of America*, 36(1):48–49, 1950.

John Nash. Non-Cooperative Games. *The Annals of Mathematics*, 54(2):286–295, 1951.

Robert Nau, Victor Richmond Jose, and Robert Winkler. Scoring Rules, Entropy, and Imprecise Probabilities. In *Proceedings of the Fifth International Symposium on Imprecise Probability: Theories and Applications*, pages 307–315, 2007.

Athanasios Papakonstantinou. *Mechanism Design for Eliciting Costly Observations in Next Generation Citizen Sensor Networks.* PhD thesis, School of Electronics and Computer Science, University of Southampton, 2010.

Athanasios Papakonstantinou, Alex Rogers, Enrico Gerding, and Nicholas Jennings. A Truthful Two-Stage Mechanism for Eliciting Probabilistic Estimates with Unknown Costs. In *Proceedings of the 18th European Conference on Artificial Intelligence (ECAI 2008), Patras, Greece*, pages 448–452, 2008.

Athanasios Papakonstantinou, Alex Rogers, Enrico Gerding, and Nicholas Jennings. Mechanism Design for Eliciting Probabilistic Estimates from Multiple Suppliers with Unknown Costs and Limited Precision. In *Proceedings of The 11th International Workshop on Agent Mediated Electronic Commerce (AMEC 2009)*, pages 111–124, 2009.

Athanasios Papakonstantinou, Alex Rogers, Enrico H Gerding, and Nicholas R Jennings. Mechanism design for the truthful elicitation of costly probabilistic estimates in distributed information systems. *Artificial Intelligence*, 175(2):648—672, 2011.

D. Pudjianto, C. Ramsay, and G. Strbac. Virtual power plant and system integration of distributed energy resources. *IET Renewable Power Generation*, 1(1):10, March 2007.

R. Tyrrell Rockafellar. *Convex Analysis.* Princeton University Press, 1970.

Saifur Rahman, Manisa Pipattanasomporn, and Yonael Teklu. Intelligent Distributed Autonomous Power Systems. In *Proceedings of Power Engineering Society General Meeting, 2007. IEEE*, pages 1–8, 2007.

Sarvapali D. Ramchurn, Perukrishnen Vytelingum, Alex Rogers, and Nicholas R. Jennings. Agent-based homeostatic control for green energy in the smart grid. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 2(4):35, July 2011.

Valentin Robu, Ramachandra Kota, Georgios Chalkiadakis, Alex Rogers, and Nicholas R. Jennings. Cooperative virtual power plant formation using scoring rules. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, pages 1165–1166, Valencia, Spain, June 2012.

Thornton B. Roby. Belief states and the uses of evidence. *Behavioral Science*, 10(3): 255–270, 1965.

Alex Rogers, Sasan Maleki, Siddhartha Ghosh, and Nicholas R Jennings. Adaptive Home Heating Control Through Gaussian Process Prediction and Mathematical Programming, May 2011.

Leonard Savage. Elicitation of Personal Probabilities and Expectations. *Journal of the American Statistical Association*, 66(336):783–801, 1971.

Leonard Savage. On Rereading R.A. Fisher. *The Annals of Statistics*, 4(3):441–500, 1976.

Mark J Schervish. A General Method for Comparing Probability Assessors. *The Annals of Statistics*, 17(4):1856—1879, 1989.

Sven Seuken, David C. Parkes, Eric Horvitz, Kamal Jain, Mary Czerwinski, and Desney Tan. Market user interface design. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC 2012)*, page 898, New York, New York, USA, June 2012. ACM Press.

Peng Shi, Vincent Conitzer, and Mingyu Guo. Prediction Mechanisms That Do Not Incentivize Undesirable Actions. In *Proceedings of the 5th International Workshop on Internet and Network Economics (WINE 2009)*, pages 89–100, Berlin, Heidelberg, 2009.

Carl-Axel S Staël von Holstein. A Family of Strictly Proper Scoring Rules Which Are Sensitive to Distance. *Journal of Applied Meteorology*, 9(3):360–364, 1970.

U.S. Department of Energy. Grid 2030: A National Vision for Electricity's Second 100 Years. Technical report, U.S. Department of Energy, Office of Electric Transmission and Distribution, 2003.

P P Varaiya, F F Wu, and J W Bialek. Smart Operation of Smart Grid: Risk-Limiting Dispatch. *Proceedings of the IEEE*, 99(1):40—57, January 2011.

William Vickrey. Counterspeculation, Auctions and Competitive Sealed Tenders. *The Journal of Finance*, 16:8–37, 1961.

Robert Winkler. The Quantification of Judgment: Some Methodological Suggestions. *Journal of the American Statistical Association*, 62(320):1105–1120, 1967.

Robert Winkler. Evaluating Probabilities: Asymmetric Scoring Rules. *Management Science*, 40(11):1395–1405, 1994.

Michael Wooldridge. *Multiagent Systems*, chapter Intelligent Agents, pages 27—77. Massachusetts Institute of Technology, 1999.

Michael Wooldridge and Nicholas R Jennings. Intelligent Agents: Theory and Practice. *The Knowledge Engineering Review*, 10:115–152, 1995.

Fredrik Ygge and Hans Akkermans. Power Load Management as a Computational Market. In *Second International Conference on Multiagent Systems*, pages 393–400, Menlo Park, California, 1996.

Fredrik Ygge, Hans Akkermans, Arne Andersson, Marko Krejic, and Eric Boertjes. The HomeBots System and Field Test: A Multi-Commodity Market for Predictive Power Load Management. In *Proceedings of the Fourth International Conference and Exhibition on The Practical Application of Intelligent Agents and Multi-Agents (PAAM'99)*, London, 1999.

Aviv Zohar and Jeffrey Rosenschein. Mechanisms for information elicitation. *Artificial Intelligence*, 172(16—17):1917–1939, 2008.