Mechanism Design for Aggregated Demand Prediction in the Smart Grid

Harry Rose and **Alex Rogers** and **Enrico H. Gerding**
Electronics and Computer Science, University of Southampton, UK

Abstract

This paper presents a novel scoring rule-based mechanism that encourages agents to produce costly estimates of future events and truthfully report them to a centre when the budget for payments to the agents is itself determined by their reports. This is applied to a model of aggregated demand prediction within a microgrid where, given estimates of future consumptions, an aggregator must optimally purchase electricity for a set of homes, each represented by self-interested, rational home agents. This in turn reduces the need for costly standby generation within the grid. The aggregator has prior information about the amount each home will consume, and determines the amount to pay each agent based on savings resulting from using the agents' reported information, over its own prior information. Agents use sensory information regarding their property and its occupants to generate these estimates, which they transmit to the aggregator using smart grid technology. The proposed mechanism is dominant strategy incentive compatible and empirical evaluation shows that it encourages agents to exert effort in producing precise estimates. We show that the mechanism is ex ante individually rational for the aggregator, and that it outperforms a simpler mechanism whereby savings are distributed evenly.

1 Introduction

The supply and demand of electricity, unlike other commodities, must be constantly balanced in real time. However, adjusting the amount of electricity being generated on the grid in real time is a costly procedure, and thus electricity markets are designed to encourage the participant to make accurate estimates of the electricity that they will generate or consume. As in the BETTA arrangements in the UK, this is typically done through a combination of forward markets, where generators and retailers exchange bilateral contracts, and a balancing market in which participants must buy or sell any real-time electricity imbalance at a cost penalty. To date, predictions of consumer demand have been made using statistical techniques applied to historical data. However, with the cost of electronics ever decreasing, the deployment of sensors and monitoring equipment within homes and buildings is becoming evermore ubiquitous. As such, a wealth of information regarding the electrical consumption and behavioural patterns of a building's occupants is becoming available for exploitation by the electricity industry. For example, non-intrusive appliance load monitoring can be used to build profiles of device use within a home such that usage patterns can be determined simply by sampling the overall consumption of the building (Zeifman and Roth 2011; Gupta, Reynolds, and Patel 2010). Moreover, sensors can be used to determine the characteristics of the building itself – its heating characteristics, and heat-leakage profiles for example. Furthermore, with the gain in popularity of cloud-services, more and more user information can be found on the web. One online source of information could be the occupants' online diaries, detailing not only periodic events that are easy to predict such as going to work, but also one-off events that have large impacts on power consumption such as social events within the house, or holidays.

It is clear then, that information is in abundance, and with the advent of new, smart grid technology, this information is becoming more accessible to novel services who can aggregate said information and use it to optimise the way electricity is purchased for those houses. This is achieved because the smart grid develops the electrical grid into one not only of energy but also data. In such a setting, autonomous *home agents* can be tasked with collecting information specific to their individual homes' consumptions, requiring minimal interaction from the home owner. Moreover, the bidirectional flow of real-time information enabled by the smart grid can be exploited to allow novel aggregation services to gather this predictive data in order to better predict and purchase electricity for their customers. This will reduce the need for inefficient standby generation; a property that makes this particularly useful in microgrid settings, as well as reducing carbon emissions and costs. Moreover, if these companies were to reward their customers for providing this information by paying them a share of the savings, both parties would benefit. Nevertheless, any such payment scheme would have to suitably reward precise and accurate estimates of a homes energy consumption, and be robust against individuals attempting to game the system.

To this end, recent research has begun to use scoring rules to elicit probabilistic estimates of future events. For example, Miller, Resnick, and Zeckhauser (2005) discuss a peer-prediction mechanism whereby agents are rewarded based on how their beliefs correspond to reports of the same event by other agents. They also discuss how such scoring rule-based mechanisms can be scaled such that agents are incentivised to exert effort in generating their reports. More recently, Papakonstantinou et al. (2011) extend that work to develop a mechanism that elicits estimates of an event with a specific precision and also discuss eliciting estimates from multiple agents about a common event and fusing those estimates to obtain a single estimate with a given precision (Papakonstantinou et al. 2011). Nau, Jose, and Winkler (2007) extend work in Winkler (1994) to discuss the application of weighted scoring rules, based on the traditional spherical and quadratic rules, that take into account not only a probabilistic estimate to be scored, but also a base estimate representing prior information. This allows agents to be scored on the *value* of the information they report, so agents reporting information to the centre that is already known will obtain a lower score than those who report more useful information.

However, none of these mechanisms address the key challenge in our setting, whereby the payment that is redistributed to the agents is dependent on the actual estimates that they submit. All previous work has assumed that this payment is constant or determined externally by the centre. Furthermore, where this work has considered multiple reports (namely Miller, Resnick, and Zeckhauser (2005) and the papers by Papakonstantinou et al.), it has done so in a setting where these reports are fused together. In contrast, in our setting, the aggregator must determine the total demand and so must convolve the individual homes' energy consumption estimates. Finally, many of these individual solutions have unresolved issues that limit their applicability. For example, the rules presented in Nau, Jose, and Winkler (2007) go partway towards addressing the issue of rewarding agents based on extra information that they might provide to the aggregator. However, in the continuous domain they are unbounded, and as such are problematic for use in mechanisms that reward agents based on those scores.

Thus, to address these shortcomings, in this paper we present a scoring rule-based mechanism that rewards agents based on the incremental value they provide to the centre in a setting where the total payments to the agents is also determined by their own reports. We apply this to a model of aggregated demand prediction whereby an aggregator collects estimates of future consumptions of houses over which it is responsible to optimally purchase electricity.

In more detail, in this paper we make the following contributions to the state of the art:

- We present a new scoring rule-based mechanism named *sum of others' plus max*, which we apply to an aggregation scenario in the smart grid that rewards agents based on the savings made by the aggregator when using the agents' estimates over the aggregator's prior estimates.
- We prove that this mechanism is dominant strategy incentive compatible and weakly budget balanced.
- We compare our mechanism using a computational approach to find the equilibrium state, to a benchmark mechanism applied to the same scenario, and show that the sum of others' plus max mechanism results in a greater social welfare while allowing the aggregator to keep a greater percentage of its savings. Specifically, we compare our mechanism to a simple mechanism that equally divides savings amongst agents.

The rest of this paper is structured as follows: Section 2 presents a model of aggregated demand prediction in the smart grid. Section 3 presents the two mechanisms described above, and discusses their properties. Section 4 empirically evaluates how agents behave when participating in the mechanisms. Finally, Section 5 concludes this paper.

2 Information Aggregation Problem

We now present a formulation of the information aggregation problem for demand prediction within the smart grid. We discuss a scenario consisting of two types of agents – a single *aggregator agent*, and N *home agents*, $i \in H$, where $H = \{1, \cdots, N\}$. The aggregator's job is to gather information about future electricity consumption of a set of homes and then buy electricity for those homes. The homes each have their own agent, whose job it is to collect specific, detailed consumption information about the home for which it is responsible and then to report an estimate of future consumption to the aggregator. The aggregator can then use this information to make better predictions of the future aggregate consumptions, which, due to the design of the electricity markets described earlier, reduces the total cost of the electricity consumed for all the homes.

In more detail, each day, T, has one time period for which the aggregator must purchase the amount of electricity it expects its agents to consume. The aggregator can purchase electricity in one of two markets, dependent on the time at which the electricity is being purchased. Electricity can be bought one day ahead of its consumption in the *forward market*, in which case it costs f per unit of electricity. At the end of each day, the aggregator is charged for any imbalance between the amount it purchased for consumption and the amount it actually consumed. We say these transactions are performed in the *balancing market* in which the prices are designed by the market regulator to penalise suppliers and consumers who do not generate or consume as they predicted. The price at which the grid buys back excess electricity, the system buy price, is $f - \delta^b$ per unit, and the cost per unit of electricity bought from the grid to fill any deficit, the system sell price, is $f + \delta^s$. Therefore, the total cost of consuming ω units of electricity when χ units are initially bought is given by:

$$\kappa\left(\omega \,|\, \chi\right) = f \cdot \chi + (\omega - \chi) \cdot \begin{cases} (f - \delta^b), & \chi > \omega \\ (f + \delta^s), & \chi < \omega \end{cases} \quad (1)$$

Each agent i has an estimate of its future consumption, x_i, represented by a Gaussian distribution[1]. Each belief is parametrised by a mean μ_i, and precision $\theta_i = 1/\sigma_i^2$, such that $x_i = \langle \mu_i, \theta_i \rangle$. Providing this estimate incurs a cost defined by $c(\alpha_i, \theta_i) = \alpha_i \theta_i$ such that more precise estimates are more costly to produce. This cost may reflect

[1]While other distributions can be applied, Gaussian distributions are used here since they well represent additive errors and are computationally attractive since the sum of two Gaussian distributions is itself a Gaussian.

the time, inconvenience or computational expense involved in generating useful estimates. The aggregator also maintains its own belief about what each agent i, will consume, $x_{a,i} = \langle \mu_{a,i}, \theta_{a,i} \rangle$ generated through conventional statistical analysis of historical consumption.

The day before the electricity is required, the aggregator asks each agent to report its estimate of tomorrow's consumption, $\hat{x}_i$, as a Gaussian with mean, $\hat{\mu}_i$, and precision, $\hat{\theta}_i$. Agents strategise over the precision of the estimate that they actually generate, θ_i, and also the precision that they report to the aggregator, $\hat{\theta}_i$. The estimates reported by the home agents are compared by the aggregator to its own information, resulting in the set of aggregated estimates $\hat{\boldsymbol{x}} = \langle x_1^*, \cdots, x_N^* \rangle$, when $x_i^* = \hat{x}_i$ if $\hat{\theta}_i > \theta_{a,i}$ and $x_i^* = x_{a,i}$, when $\hat{\theta}_i \leq \theta_{a,i}$. In order to make our notation less verbose, we here define the aggregated mean of all estimates in $\hat{\boldsymbol{x}}$ as $\hat{\mu} = \sum_{x_i^* \in \hat{\boldsymbol{x}}} \mu_i^*$, and the aggregated precision as $\hat{\theta} = 1/(\sum_{x_i^* \in \hat{\boldsymbol{x}}} 1/\theta_i^*)$. Similarly, for the aggregator's beliefs, $\boldsymbol{x}_a = \langle x_{a,1}, \cdots, x_{a,N} \rangle$.

It is important to note that the aggregator does *not* fuse any estimates reported by the agents as it is unaware of any treatment of the information by the home agents. For example, it cannot fuse its own historical estimates with the estimates reported by the agents as the home agents may have already used the same historical data in generating their estimates, and thus, these estimates would not be independent. Thus, the aggregator acts conservatively, assumes that its own and the agents' estimates are correlated and uses the one with greater precision in the convolution. Also, note that, as the aggregated belief is the convolution of Gaussian distributions reported by the agents, the result is again a Gaussian.

The aggregator then performs the following optimisation in order to determine how much electricity it must buy in the forward market such that its total expected cost is minimised given an estimated aggregate consumption of $\hat{\boldsymbol{x}}$.

$$\begin{aligned}\chi(\hat{\boldsymbol{x}}) = \underset{z \in \Omega}{\operatorname{argmin}}\ & f \cdot z - \\ & \int_0^z (z-y)(f-\delta^b)\mathcal{N}\left(y; \hat{\mu}, \hat{\theta}\right) \mathrm{d}y + \\ & \int_z^\infty (y-z)(f+\delta^s)\mathcal{N}\left(y; \hat{\mu}, \hat{\theta}\right) \mathrm{d}y\end{aligned} \tag{2}$$

That is, it finds an amount to buy that minimises the cost of buying z units at f in the forward market, and minimises the total cost of the expected error in the balancing market.

At the end of each day, the actual amount consumed by each agent, $\boldsymbol{\omega} = \langle \omega_1, \cdots, \omega_N \rangle$, is known by the aggregator. The total consumption is defined as $\omega = \Sigma_{\omega_i \in \boldsymbol{\omega}} \omega_i$. Once the consumption of each agent has been realised the aggregator is then able to calculate precisely how much it saved by using the agents information over its own. Using the cost function, κ from Equation 1, and the optimisation function from Equation 2 the saving made by the aggregator using the agents estimates $\hat{\boldsymbol{x}}$, over its own estimates $\boldsymbol{x}_a$, when the agents consume $\boldsymbol{\omega}$ is defined as:

$$\Delta(\hat{\boldsymbol{x}}, \boldsymbol{x}_a, \boldsymbol{\omega}) = \kappa(\boldsymbol{\omega} \mid \chi(\boldsymbol{x}_a)) - \kappa(\boldsymbol{\omega} \mid \chi(\hat{\boldsymbol{x}})) \tag{3}$$

Similarly, the savings made by all agents *except* agent i – those made by the aggregator had agent i not participated in the mechanism – is defined as above except using $\hat{\boldsymbol{x}}_{-i} = \langle x_1^*, \cdots, x_{i-1}^*, x_{i+1}^*, \cdots, x_N^* \rangle$, $\boldsymbol{x}_{a,-i} = \langle x_{a,1}, \cdots, x_{a,i-1}, x_{a,i+1}, \cdots, x_{a,N} \rangle$ and $\boldsymbol{\omega}_{-i} = \langle \omega_1, \cdots, \omega_{i-1}, \omega_{i+1}, \cdots, \omega_N \rangle$, in place of $\hat{\boldsymbol{x}}$, $\boldsymbol{x}_a$ and $\boldsymbol{\omega}$ respectively.

The aggregator then allocates a fraction λ of the savings earned to the reward mechanism that distributes it to the agents. Each agent i is then rewarded for its information by an amount defined by some reward function over the reported means, variances, and outcomes $P(\hat{\boldsymbol{x}}, \boldsymbol{\omega})$. Therefore, the agent's utility is given by:

$$U(\hat{\boldsymbol{x}}, \boldsymbol{\omega}, \alpha_i) = P(\hat{\boldsymbol{x}}, \boldsymbol{\omega}) - c_i(\alpha_i, \theta_i) \tag{4}$$

Since the aggregator uses the estimates provided by the agents to purchase an amount of electricity that minimises its expected costs, it must encourage the agents to truthfully report estimates that are both accurate and precise. We refer to an accurate estimate as having a mean close to the realised outcome, and a precise estimate as one with high precision. It does this by carefully designing the payments it gives to agents for their information such that those payments are maximised in expectation when the agents behave as required. In the next section, we use mechanism design to design such payment schemes.

3 Mechanisms

We now present two mechanisms that allocate rewards to agents for their information. A mechanism specifies a *transfer function*, which defines the reward an agent receives for a given reported estimate, $\hat{x}_i$, when an outcome, ω_i, is realised. Specifically, we consider three properties that are desirable in the scenario presented earlier. First, our mechanism should exhibit *individual rationality*. That is, all agents gain a positive utility from participating in the mechanism in expectation. This is an essential requirement of any mechanism designed for use within an aggregation service to which customers may opt out – people will simply not use the service if they expect to be worse off by so doing. Second, our mechanism should be *incentive compatible*, which means that an agent maximises its expected utility by truthfully reporting its estimate. This has obvious advantages in the aggregation scenario described earlier – the aggregator needs to know the real estimates the agents hold in order to generate accurate estimates of their aggregate future consumption. Third, it should be *weakly budget balanced*, which states that the aggregator does not run into deficit after paying the agents for their estimates. We consider a mechanism to be budget balanced if the aggregator spends in total equal or less than it would have, had it only used its own estimates and not elicited estimates from the home agents. This is important because it is the aggregator that will implement this mechanism, and if he does not expect to profit from it, he will simply choose another model.

Expanding on the second requirement (incentive compatibility), it is important to note that the ability to strategise over whether or not to misreport is not unique to the home agents. If individual *home owners* were to find a way to profit by misreporting the information they supply to the home agent, it would be in their best interest to do so. Thus an incentive compatible mechanism results in both home agents and home owners reporting information truthfully.

Given this, we first discuss a simple mechanism whereby the savings made by the aggregator are equally divided amongst the home agents. Then, we discuss a further mechanism which uses the spherical scoring rule to define the proportion of the savings distributed to each agent. Table 1 presents the mechanisms and their differing properties.

3.1 Uniform Scoring

The first mechanism presented in this paper simply divides the savings made by the aggregator equally amongst the agents. In this case, the reward given to each agent is:

$$P^{\mathbf{U}}\left(\hat{\boldsymbol{x}}, \boldsymbol{x}_a, \boldsymbol{\omega}, N\right) = \frac{\lambda \cdot \Delta\left(\hat{\boldsymbol{x}}, \boldsymbol{x}_a, \boldsymbol{\omega}\right)}{N}$$

Theorem 3.1 *In the uniform scoring mechanism, truth telling is a Nash equilibrium.*

Proof The aggregator buys an amount of electricity for the agents that minimises the total expected cost based on the agents' reported estimate. Clearly if all agents report truthfully, one agent deviating will only cause a larger error between the amount consumed and purchased, resulting in less savings to be distributed to the agents and therefore a lower utility for that agent. Therefore, when all agents report truthfully, a single agent is unable to improve its expected utility by misreporting. However, truth telling is not dominant. If an agent knows its neighbour will misreport, it can obtain a better expected reward by also misreporting such that the two errors cancel each other out. □

This mechanism will provide us with a benchmark on which to base our comparisons. This simple mechanism is both incentive compatible and budget balanced when used in this model. However, it does not provide good fairness properties. Using this rule, all agents are rewarded equally irrespective of their actual contribution. An ideal mechanism would reward the agents whose estimates made the most significant increase in the aggregator's savings with a greater payment. Furthermore, the fact that truth telling is only a Nash equilibrium means that agents can potentially expect to benefit from misreporting their estimates if they believe the other agents will also misreport. A stronger mechanism would be strict dominance incentive compatible. That is, an agent's utility is maximised when reporting truthfully *regardless* of its belief of the other agents' actions. With this in mind, we propose the sum of others' plus max mechanism, based on strictly proper scoring rules. Strictly proper scoring rules are functions that are maximised in expectation when an agent reports its estimates truthfully. Next, we present next some background theory on the *spherical* scoring rule, and then the *sum of others' plus max* mechanism that uses it to achieve strict dominance incentive compatibility.

3.2 Spherical Scoring Rule

The mechanism in the next section is based on the spherical scoring rule. For a given prediction of an event with mean $\hat{\mu}_i$, and precision $\hat{\theta}_i$, and a realisation of that event, ω, the spherical rule is defined as follows:

$$S\left(\omega_i \,\middle|\, \hat{\mu}_i, \hat{\theta}_i\right) = \frac{\mathcal{N}\left(\omega; \hat{\mu}_i, \hat{\theta}_i\right)}{\sqrt{\int_{-\infty}^{\infty} \mathcal{N}\left(x; \hat{\mu}_i, \hat{\theta}_i\right)^2 \mathrm{d}x}} \quad (5)$$

The spherical rule is one of three *strictly proper* scoring rules often studied in literature – the other two being the logarithmic, and quadratic scoring rules. The term strictly proper means that the score awarded by the function is maximised exclusively when the agent truthfully reports its estimate. The spherical rule was chosen over the much simpler logarithmic rule because it has a strict lower bound of 0, whereas the logarithmic rule is unbounded. As a result, the use of the logarithmic scoring rule could theoretically end with a customer becoming forever in debt to the aggregation company after having received a score of $-\infty$. This is clearly unsatisfactory, and it could be argued that this fact, no matter how rare its occurrence, could dissuade users from ever joining the aggregation service.

3.3 Sum of Others' plus Max

The sum of others' plus max mechanism improves upon the uniform mechanism by rewarding each agent individually dependent on their reports. The result is that agents are more fairly rewarded as those who provide precise and accurate estimates receive a higher payment from the aggregator than those who do not. This mechanism takes into account not only the spherical score achieved by the agent, but also the spherical scores achieved by the other agents in the system. Crucially, payments are then determined by multiplying the ratio of those scores by the savings made by *other* agents in the system, which is necessary in order to preserve the incentive compatibility of the mechanism.

Given all agents' estimates, agent i is paid:

$$P^{\mathbf{S}}\left(\hat{\boldsymbol{x}}, \boldsymbol{x}_a, \boldsymbol{\omega}\right) = \frac{S\left(\omega_i; \hat{\mu}_i, \hat{\theta}_i\right) \cdot \lambda \cdot \Delta\left(\hat{\boldsymbol{x}}_{-i}, \boldsymbol{x}_{a,-i}, \boldsymbol{\omega}\right)}{S\left(\omega_j; \omega_j, \bar{\theta}\right) + \Sigma_{x_j \in \hat{\boldsymbol{x}}_{-i}} S\left(\omega_j; \hat{\mu}_j, \hat{\theta}_j\right)}$$

Here the $S\left(\omega_j; \omega_j, \bar{\theta}\right)$ term represents the maximum score that can be achieved by an agent – the score achieved when being accurate (reporting an estimate with mean ω_j and ω_j occurring) and with the maximum precision of $\bar{\theta}$. The payments awarded to agents by this mechanism are fairer than those of the uniform scoring mechanism insofar as the agents are directly compared with each other by dividing each agent's score by the sum of all other agents' scores. This is similar to calculating the percentage of the total score of all agents that agent i contributed. However, note that in this calculation, agent i's score is *not* included in the divisor, and nor is the savings through using his information included in the amount to be distributed to him. This is necessary to maintain incentive compatibility in the mechanism. Moreover, it is essential to divide the prescaled score by the sum of the other agents' prescaled scores *plus the maximum score* in order to maintain budget balance (consider the case in which agent i scores more highly than agent j).

Name	Reward Function	Truth Telling	Budget Balance
Uniform Scoring	$P^{\mathbf{U}}\left(\hat{\boldsymbol{x}}, \boldsymbol{x}_a, \boldsymbol{\omega}, N\right)=\frac{\lambda \cdot \Delta(\hat{\boldsymbol{x}}, \boldsymbol{x}_a, \boldsymbol{\omega})}{N}$	Nash equilibrium	Strong
Sum of Others' plus Max	$P^{\mathbf{S}}\left(\hat{\boldsymbol{x}}, \boldsymbol{x}_a, \boldsymbol{\omega}, \bar{\theta}\right)=\frac{S\left(\omega_i;\hat{\mu}_i,\hat{\theta}_i\right)\cdot\lambda\cdot\Delta(\hat{\boldsymbol{x}}_{-i},\boldsymbol{x}_{a,-i},\boldsymbol{\omega})}{S\left(\omega_j;\omega_j,\bar{\theta}\right)+\Sigma_{x_j\in\hat{\boldsymbol{x}}_{-i}}S\left(\omega_j;\hat{\mu}_j,\hat{\theta}_j\right)}$	Strictly dominant	Weak

Table 1: Comparison of the two mechanisms discussed in this paper

Theorem 3.2 *In the sum of others' plus max mechanism, truth telling is strictly dominant.*

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 that, the sum of others' plus max rule 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 is a unique maximum when 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. □

There are several advantages to using the sum of others' plus max mechanism over the simple uniform one presented earlier. Firstly, truth telling strictly dominates all other strategies, as such agents need not compute their utilities over the whole strategy space but can simply report their true estimates. Moreover, each agent will maximise his expected utility by truthfully reporting his estimate to the aggregator *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 one agent were to learn that its neighbour were 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 interim individually rational for the aggregator – a weaker concept than the ex post individual rationality exhibited by the uniform mechanism. That is, in the sum of others' plus max mechanism, when the home agents are reporting truthfully the aggregator will in expectation make a profit. However, individual instances may occur in which the aggregator makes a loss through using the agents' information. The sum of others' plus max mechanism will also rarely distribute 100% of the allocated savings to the home agents – the total amount it distributes is dependent on each agents' score. In these instances, any undistributed savings are simply kept by the aggregator.

4 Empirical Evaluation of Mechanisms

In generating their estimates, the agents incur some cost that is proportional to the precision of the generated estimate. Agents can strategise over both the *actual* precision of their estimate, θ_i, and the precision they *claim* their estimate to be when reporting it to the aggregator, $\hat{\theta}_i$. Agents can also strategise over misreporting $\hat{\mu}_i$. However, as shown in the previous section, the uniform and sum of others' plus max mechanisms are incentive compatible and therefore rational agents will always truthfully report their estimates ($\hat{x}_i = x_i$). Nevertheless, the actual precision of the estimates the agents generate will determine their expected utility, and the choice of precision is dependent on the choices of the other agents. Thus, in this section we calculate the equilibrium choice of which precision to generate made by the agents.

To achieve this, we simulate a microgrid consisting of 100 houses (each with one home agent) with cost coefficients, α_i distributed uniformly between 0.5 and 1.0; and an aggregator agent. We then use iterated best response to find the equilibrium strategy of the agents. To do so, each agent chooses a strategy (precision) that maximises its utility given the strategies of the other agents in the previous iteration. In the first game, the agents' knowledge of the other agents' strategies is initialised to $\bar{\theta} = 10$. Note that this does not describe the real setting, but rather, the computational algorithm we use to find this equilibrium point. In order to aid convergence to a single, pure strategy as opposed to a cycle, each agent takes turns to set its strategy as in the alternate-move Cournot dynamic (Fudenberg and Levine 1998, pp.10). Each agent assumes that the other agents' strategies are fixed and that the agents will always report the same precision as they had done in the previous game. The agent then chooses its precision by analytically maximising the agent's expected utility function, given by:

$$\bar{U}\left(\hat{\theta}_i, \tilde{\theta}_{-i}, \theta_{a,i}, \theta_{a,-i}\right) = \iint_{-\infty}^{\infty} P\left(\hat{\boldsymbol{x}}, \boldsymbol{x}_a, \boldsymbol{\omega}\right) - c(\alpha_i, \theta_i) \cdot \mathcal{N}\left(\omega_i; \mu_i, \theta_i\right) \cdot \mathcal{N}\left(\omega_{-i}; \mu_{-i}, \tilde{\theta}_{-i}\right) \mathrm{d}\omega_i \, \mathrm{d}\omega_{-i}$$

where P is a reward function defined in Section 3. Similar experiments were performed with a range of numbers of houses, N. These yielded similar results, although the benefit of the sum of others' plus max mechanism over the uniform mechanism reduces as N decreases.

The consumption, ω_i of each agent, i, for each trial is uniformly, randomly distributed over the range $[30, 50]$. In each trial of every iteration, the precision, θ_i, of the estimate the agent generates is determined as above, and the mean, μ_i, drawn from the Gaussian distribution by $\mu_i \sim \mathcal{N}\left(\omega_i, \theta_i\right)$.

Figure 1 shows the mean of the sum of utilities attained by the home agents for simulations whereby the aggregator chooses to allocate a fraction $\lambda \in \{0.0, 0.1, 0.2, \cdots, 1.0\}$ of the savings made to be used to repay the agents for their information. Also plotted is the maximum possible utility for any mechanism as the agents' costs tend to zero and their precision tends to infinity. This is calculated by dividing the maximum savings the aggregator can make – the

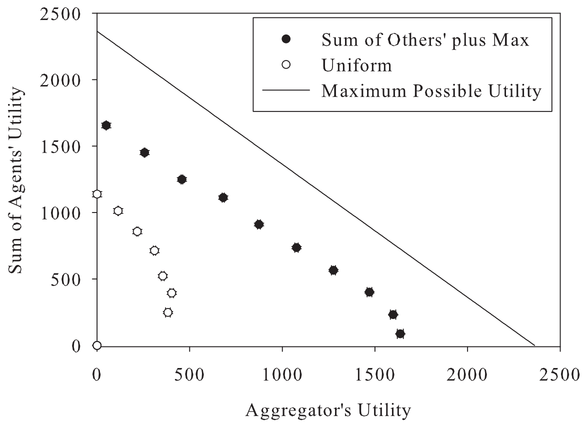

Figure 1: The sum of the home agents' utilities plotted against the aggregator's utility.

cost of electricity the aggregator would have incurred had it used its own belief minus the cost of electricity had it bought the exact correct consumption in the forward market. These savings are then divided between the aggregator and home agents' in accordance with λ. This gives us an upper bound, and a method by which to compare the two mechanisms.

We can see from Figure 1 that the sum of others' plus max mechanism is more efficient than the uniform mechanism insofar as it is closer to the upper bound on utility than the uniform mechanism. This can be explained using the shape of the agents' utility functions under both mechanisms. When strategising over the precision with which to generate a report, the peak expected utility is at a higher precision in the sum of others' plus max case than in the uniform case. This results in agents choosing to generate reports of higher precision, therefore resulting in the aggregator making greater savings from their reports. Figure 2 further confirms that the sum of others' plus max mechanism is incentivising agents to choose to generate reports of a higher precision. It can be seen that, for 100 agents, the uniform mechanism provides little incentive for agents to increase the precision of their reports. In this case, the increase in savings from an agent providing a report of higher precision, when divided equally among the 100 agents, is outweighed by the extra cost incurred by that agent. However, in the sum of others' plus max mechanism, when generating a report of higher precision, the agent's expected payment is not only increased due to the increased expected savings, but also it's increased expected score. The error bars in Figure 2 – showing standard error – also show that the fraction of the savings made by the aggregator that is paid to agents is variable. This results from the use of the spherical score in calculating the fraction of the savings to allocate to agents and the fact that agents' beliefs are not 100% accurate, and have finite precision.

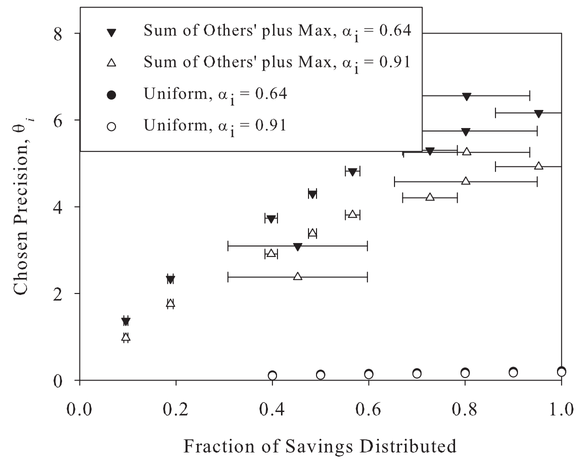

Figure 2: The precision of the estimates made by the agents for the uniform and sum of others' plus max mechanisms.

5 Conclusions

This paper discussed mechanism design for information aggregation for aggregated demand prediction in the smart grid. A model of aggregated demand prediction was presented in which an aggregator rewards a group of home agents for probabilistic estimates of their future consumption. A novel mechanism named 'sum of others' plus max' was presented, which employs the spherical scoring rule to achieve incentive compatibility. It was shown that truth telling in this mechanism is strictly dominant, as agents uniquely maximise their expected reward when truthfully reporting, *regardless* of the other agents' actions. This was compared to a simpler mechanism whereby the savings are equally divided amongst agents, and was shown to result in a greater social welfare and a greater utility for the aggregator.

Future work will investigate the use of estimates from a subset of the home agents in order to predict the consumptions of the missing agents when the aggregator cannot afford to buy reports from all agents.

References

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

Gupta, S.; Reynolds, M. S.; and Patel, S. N. 2010. Electrisense: single-point sensing using emi for electrical event detection and classification in the home. In *Proceedings of the 12th ACM international conference on Ubiquitous computing*, Ubicomp '10, 139–148. New York, NY, USA: ACM.

Miller, N.; Resnick, P.; and Zeckhauser, R. 2005. Eliciting informative feedback: The peer-prediction method. *Management Science* 51(9):1359—1373.

Nau, R.; Jose, V. R.; and Winkler, R. 2007. Scoring rules, entropy, and imprecise probabilities. In *Proceedings of the Fifth International Symposium on Imprecise Probability: Theories and Applications*, 307—315.

Papakonstantinou, A.; Rogers, A.; Gerding, E. H.; and Jennings, N. R. 2011. Mechanism design for the truthful elicitation of costly probabilistic estimates in distributed information systems. *Artificial Intelligence* 175(2):648 – 672.

Winkler, R. 1994. Evaluating probabilities: Asymmetric scoring rules. *Management Science* 40(11):1395—1405.

Zeifman, M., and Roth, K. 2011. Nonintrusive appliance load monitoring: Review and outlook. *Consumer Electronics, IEEE Transactions on* 57(1):76 –84.