A Scoring Rule-Based Mechanism for Aggregate Demand Prediction in the Smart Grid

Harry Rose, Alex Rogers, Enrico H. Gerding
Electronics and Computer Science, University of Southampton
{htrose,acr,eg}@ecs.soton.ac.uk

ABSTRACT

This paper presents a novel scoring rule-based strictly dominant incentive compatible mechanism that encourages agents to produce costly estimates of future events and report them truthfully to a centre. Whereas prior work has assumed a fixed budget for payment towards agents, this work makes use of prior information held by the centre and assumes a budget that is determined by the savings made through the use of the agents' information over the centre's own prior information. This mechanism is compared to a simple benchmark mechanism wherein the savings are divided equally among all home agents, and a cooperative solution wherein agents act to maximise social welfare. Empirical analysis is performed in which the mechanism is applied to a simulation of the smart grid whereby an aggregator agent must use home agents' information to optimally purchase electricity. It is shown that this mechanism achieves up to 77% of the social welfare achieved by the cooperative solution.

Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]: Multiagent systems

General Terms

Economics, Theory

Keywords

Mechanism Design, Smart Grid, Information Aggregation, Scoring Rules

1. INTRODUCTION

There are numerous scenarios in which, in order for a centre to optimally perform a task, it must gather predictive information from other, potentially non-cooperative experts such that the centre can optimally plan for future events. In such situations, the centre often incurs a cost related to the imprecision of its predictions, and consequently, the ability to produce precise estimates results in a reduction of these costs. In non-cooperative situations, it becomes necessary for the centre to make payments to the experts in order to encourage them to report any relevant information that they have. When these experts are self-interested rational agents, the payments must be carefully designed in order to elicit the behaviour that the centre requires.

As an example, consider the instance of the above setting on which this paper focuses; information aggregation within the smart grid. In this scenario, an *aggregator agent* must purchase electricity for a set of homes, each of which is represented by its own *home agent*. Based on historical evidence, the aggregator has some belief of what each house will consume at a given future time, but each house has a wealth of information within it that can, *at a cost*, be collected and processed by the home agent in order to be used to make more precise estimates. It is the home agents' jobs to gather this information on behalf of the home owners and transmit it to the aggregator agent in the form of a probability distribution over the houses' possible consumptions at the specified time. The aggregator must pay a penalty for any difference between the amount it purchases and the amount its customers consume. Thus, with more precise information, the aggregator will in expectation make some savings. The aggregator then distributes a portion of these savings to the home agents as a reward for their information.

The truthful elicitation of information from self-interested agents has already been the subject of much attention in the literature. It has been shown that in scenarios where a number of expert agents can be called upon to make probabilistic estimates of some value, peer prediction techniques can be employed to ensure those agents report truthfully [4, 6]. Peer prediction rewards experts by fusing reports from a reference agent and the agent being scored; both of whom are making estimates of the same variable. After the value of the variable becomes known, the agent whose report is being evaluated is rewarded according to how its report affected the estimated likelihood of the realised event. This evaluation makes use of *proper scoring rules* – functions of probabilistic reports and outcomes that return a score that is only maximised in expectation when the agent reports truthfully. Indeed, scoring rules have also been the focus of much attention, since their original use in meteorology [1], through more general applications of evaluating predictors [8], and their more recent application in the field of prediction markets [7] wherein experts are asked by a centre to give a probabilistic response to some question, and also in the fair division of rewards among agents performing a task [2]. Moreover, strictly proper scoring rules have been generalised to take into account a prior distribution representing knowledge held by the centre [5].

However, there are four main limitations to these works

Appears in: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, Conitzer, Winikoff, Padgham, and van der Hoek (eds.), 4-8 June 2012, Valencia, Spain.

that limit their applicability to the problem discussed in this paper. Firstly, the works assume there exists multiple experts from which the centre can elicit information. In doing this, the experts' reports can be compared to each other, and the experts' payments based upon how correlated theirs and the other agents' reports were. This results in payments that are *Nash incentive compatible*, whereas a preferable solution concept would be *dominant strategy incentive compatible*, whereby truth telling is a dominant strategy regardless of the other agents' behaviour. Furthermore, in the scenario above, only a single agent exists per home, and agents are unable to measure their neighbours' demand. Consequently, agents' reports cannot be directly compared as the events they are predicting are unique to themselves. Secondly, in prior work, the reports from the experts are fused with one-another, whereas in the scenario above, the centre is interested in the *cumulative* demand of agents, and therefore must take the convolution of the reports. Thirdly, the solutions do not take into account prior knowledge held by the centre, with the exception of [5]. In the scenario above, the centre has prior information about each home, and will make savings based upon the precision of the information he uses. The mechanism should take this into account such that the aggregator does not pay for information less precise than his own, and does not run a deficit in paying for said information. The scoring rules in [5] have been adapted to take into account prior information. However, the computation of the scores can be problematic and are unbounded in the continuous domain. Finally, the solutions assume that the budget for payment to the agents is somehow fixed, whereas we propose to make use of the fact the centre has prior information, and base agents' rewards on the *savings made by the centre* in using the agents' reports over his own prior information.

Against this background, we develop a novel, scoring rule-based mechanism named the *sum of others' plus max* mechanism, which distributes payments to agents from a budget that is determined by their own reports in a way that is incentive compatible, and *ex ante* weakly budget balanced (i.e. the expected sum of payments to the agents is less than or equal to the budget allocated for rewards by the centre).

In more detail, this paper makes the following contributions to the state of the art:

- We present a new scoring rule-based mechanism named *sum of others' plus max* (SOM), which we apply to the scenario of information aggregation in the smart grid. The mechanism rewards agents *using a budget determined by their own reports*, and takes into account the agents' reports *and the centre's prior information.*

- We prove this mechanism to be *dominant strategy incentive compatible* and *ex ante weakly budget balanced.*

- We compare this mechanism using a computational approach to find equilibrium states to a benchmark mechanism in which rewards are divided uniformly between agents, and the cooperative social-welfare maximising solution. In doing so, we show that this mechanism is much more efficient than the benchmark, and obtains a social welfare that is up to 77% of that of the optimal cooperative solution.

- We show that SOM reduces the risk of the aggregator making a loss compared to the uniform mechanism.

The remainder of the paper is structured as follows: Section 2 presents a formal model of the information aggregation problem applied to the smart grid. Section 3 discusses two mechanisms to reward agents and their theoretical properties – the uniform mechanism, and the sum of others plus max mechanism. Section 4 then discusses a *cooperative* solution whereby all agents try to maximise the social welfare of the system. Section 5 uses empirical analysis to compare the two mechanisms and the social welfare solution. Finally, the paper concludes in Section 6.

2. THE INFORMATION AGGREGATION PROBLEM

This section presents a formulation of the information aggregation problem for demand prediction within the smart grid. In this scenario there are two types of agents – a single *aggregator agent*, and n *home agents*, $i \in N$, where $N = \{1, \cdots, n\}$. The aggregator's job is to gather information about the 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 it 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, reduces the total cost of the electricity consumed for all the homes.

In more detail, each day, D, is divided into a number of time periods, which, for simplicity and without loss of generality, we assume to be one. For each time period, 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 of purchase. 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 in the forward market 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:

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

Each agent, i, can generate *at a cost* an estimate, x_i, of its future consumption, represented by a Gaussian distribution[1] with mean, μ_i, and precision, $\theta_i = 1/\sigma_i^2$, such that $x_i = \langle \mu_i, \theta_i \rangle$. However, it is important to note that the mechanisms generalise to any probability distribution. The cost incurred by agent i when generating an estimate of pre-

[1]Gaussian distributions were chosen as the estimates produced by the agents are likely to be dependent on numerous noisy sources of information, and by the central limit theorem, the sum of noisy data results in a Gaussian error.

cision θ_i is:

$$C(\alpha_i, \theta_i) = \alpha_i \cdot \theta_i \tag{2}$$

where α_i is some positive, real-valued constant. The aggregator also maintains its own belief about what each agent i, will consume, $x_{a,i} = \langle \mu_{a,i}, \theta_{a,i} \rangle$. However, the aggregator does not incur a cost in maintaining its belief.

The day before the electricity is required, the aggregator asks each agent to report its estimate of tomorrow's consumption, $\hat{x}_i = \langle \hat{\mu}_i, \hat{\theta}_i \rangle$, as a Gaussian distribution with mean, $\hat{\mu}_i$, and precision, $\hat{\theta}_i$. The home agents are assumed to be strategic and they will try to maximise the benefit they receive; defined broadly as some payment for their report minus the cost of generating that report. As such, agents strategise over the precision of the estimate that they actually generate, $x_i = \langle \mu_i, \theta_i \rangle$, and also the mean and precision that they *report* to the aggregator, $\hat{x}_i = \langle \hat{\mu}_i, \hat{\theta}_i \rangle$. Consequently, an agent will misreport (i.e. $\hat{x}_i \neq x_i$) if it believes doing so will gain it a greater utility.

Once received, the estimates reported by the home agents are compared by the aggregator to its own information. The aggregator does not know the correlation between the reports of the agents and its own. Thus it takes the conservative action of assuming they are perfectly correlated, and simply takes the most precise estimate of its own and the agent's. In so doing, the aggregator produces the following aggregate belief vector:

$$\boldsymbol{x} = \langle x_1^*, \cdots, x_n^* \rangle \tag{3}$$

where,

$$x_i^* = \begin{cases} \hat{x}_i, & \text{if } \hat{\theta}_i > \theta_{a,i} \\ x_{a,i}, & \text{otherwise} \end{cases} \tag{4}$$

The result of Equation 4 is that the aggregator will only use the home agent's reported estimate, $\hat{x}_i$, if said estimate is more precise than the estimate the aggregator already has, $x_{a,i}$. Therefore, the aggregator will never use information that is less precise than its own belief. In certain situations, this can result in the aggregator losing information. However, in general this behaviour is a necessary consequence of the aggregator being unaware of the correlation of the information sources being received, although domain-specific knowledge might be applied to overcome this limitation. Since the aggregator is interested in the *cumulative* predicted demand of all homes, it convolves $\boldsymbol{x}$ to calculate a distribution that represents the expected total demand.

In order to make notation less verbose, let the expected total demand according to $\boldsymbol{x}$ be $\mu = \sum_{x_i^* \in \boldsymbol{x}} \mu_i^*$, and its precision $\theta = 1/(\sum_{x_i^* \in \boldsymbol{x}} 1/\theta_i^*)$. Similarly, for the aggregator's beliefs, $\boldsymbol{x}_a = \langle x_{a,1}, \cdots, x_{a,n} \rangle$, let the mean and precision be $\mu_a = \sum_{x_{a,i} \in \boldsymbol{x}_a} \mu_{a,i}$ and $\theta_a = 1/(\sum_{x_{a,i} \in \boldsymbol{x}} 1/\theta_{a,i})$.

Once the aggregator has collected estimates from all agents, it performs an optimisation to determine the amount of electricity it must purchase in the forward market such that its total expected cost is minimised. Essentially, the aggregator tries to minimise its expected loss in the balancing markets, which it does by solving the following equation:

$$\chi(\boldsymbol{x}) = \underset{z \in \Omega}{\arg\min}\ f \cdot z - \int_0^z (z-y)(f-\delta^b)\mathcal{N}(y;\mu,\theta)\ \mathrm{d}y + \int_z^\infty (y-z)(f+\delta^s)\mathcal{N}(y;\mu,\theta)\ \mathrm{d}y \tag{5}$$

At the end of each day, the actual amount consumed by each agent, defined by $\boldsymbol{\omega} = \langle \omega_1, \cdots, \omega_n \rangle$, becomes known to the aggregator. The total consumption is defined as $\omega = \sum_{\omega_i \in \boldsymbol{\omega}} \omega_i$. Each agent then pays the aggregator for the electricity their home consumed, at a rate of f_r per unit. The aggregator can also calculate the total cost it incurred through utilising the agents' estimates, $\kappa(\omega \,|\chi(\boldsymbol{x}))$, and the cost it *would* have incurred had it simply used its own prior information, $\kappa(\omega \,|\chi(\boldsymbol{x}_a))$.

The aggregator must then decide an amount to pay each home agent, P_i. Agent i's utility is then defined as follows:

$$U_i(\boldsymbol{x}, \boldsymbol{x}_a, \boldsymbol{\omega}) = P_i(\boldsymbol{x}, \boldsymbol{x}_a, \boldsymbol{\omega}) - C(\alpha_i, \theta_i) - \omega_i \cdot f_r$$

and the aggregator's utility is defined as:

$$U_a(\boldsymbol{x}, \boldsymbol{x}_a, \boldsymbol{\omega}) = \omega \cdot f_r - \kappa(\omega|\chi(\boldsymbol{x})) - \sum_{i \in N} P_i(\boldsymbol{x}, \boldsymbol{x}_a, \boldsymbol{\omega}) \tag{6}$$

Given that the agents are *rational*, they are able to strategise over their space of reports in order to determine how precisely that they generate their estimate and then, whether or not to truthfully report that estimate to the aggregator. Indeed, even after an agent has paid to produce an estimate, it might still gain a greater reward by misreporting. However, the aggregator is able to incentivise agents to behave in certain ways by carefully designing the reward function, P_i. In so doing, it wants to achieve two main goals: to incentivise home agents to make precise estimates and to incentivise agents to report those estimates truthfully. This is a problem that is discussed in the next section.

3. NON-COOPERATIVE MECHANISMS

This section presents 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, the mechanism should exhibit *individual rationality*. That is, in expectation, all agents gain a positive utility from participating in the mechanism. 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, the 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 an accurate estimate of their aggregate future consumption. Third, it should be *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 equal to or less than it would have, had it only used its own estimates and not elicited estimates from the home agents.

Given this, the next section discusses how the savings are calculated in order to determine the aggregator's budget. Afterwards, the mechanisms that distribute this budget are discussed. First to be discussed is the *uniform mechanism*; a simple mechanism whereby the savings made by the aggregator are equally divided amongst the home agents. Next, a further mechanism named *sum of others' plus max* is discussed, which uses the spherical scoring rule in order to define the proportion of the savings distributed to each agent.

3.1 Calculation of Savings

The budget the aggregator uses to reward the home agents in the mechanisms presented here is a function of the total savings made by the aggregator buying electricity using the home agents' information over its own (if the home agents' information is more precise than the aggregator's). Formally, when the agents consume $\boldsymbol{\omega}$, their aggregated reports are $\boldsymbol{x}$, and the aggregator's aggregated prior information is $\boldsymbol{x}_a$, the savings made by the aggregator are:

$$\Delta\left(\boldsymbol{x}, \boldsymbol{x}_a, \boldsymbol{\omega}\right) = \kappa\left(\omega \,|\chi\left(\boldsymbol{x}_a\right)\right) - \kappa\left(\omega \,|\chi\left(\boldsymbol{x}\right)\right) \tag{7}$$

The aggregator may not necessarily decide to allocate the whole amount of savings to the agents' reward budget. Instead, it allocates a fraction $0 \leq \lambda \leq 1$ to distribute, thereby guaranteeing the aggregator a certain fraction of the savings. However, the aggregator is still left with the problem of allocating the λ savings to the agents such they are incentivised to report precise estimates truthfully. The next sections are devoted to describing two mechanisms. First, their formal properties are discussed, then further examination of their properties is performed using empirical evaluation.

3.2 Uniform Mechanism

The simplest 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:

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

It is clear to see that the uniform mechanism is budget balanced (i.e. it always distributes 100% of it's allocated budget) – the budget is equally split into n amounts, which are then awarded to n agents, thus always distributing 100% of the allocated budget. Further to this, Theorem 3.1 provides a proof that shows the uniform mechanism to be Nash incentive compatible (that is, reporting truthfully is a Nash equilibrium). Thus, under this mechanism, when all agents are truthful, no single agent has incentive to misreport.

THEOREM 3.1. *The uniform mechanism is Nash incentive compatible, i.e. 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, note that truth telling is not a dominant strategy; if an agent knows its neighbour will misreport, the agent can obtain a better expected reward by also misreporting such that its error cancels out the error made by the neighbour. □

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. 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 other agents will do the same. Therefore, a better solution is a mechanism that is dominant strategy incentive compatible. That is, a mechanism where an agent's utility is maximised when reporting truthfully *regardless* of its belief of the other agents' actions. With this in mind, we discuss next the *sum of others' plus max* mechanism (SOM), which uses strictly proper scoring rules in order to achieve dominant strategy incentive compatibility. Strictly proper scoring rules are functions that take a probabilistic estimate reported by an agent, and an outcome. Their expected value is maximised only when an agent truthfully reports their estimates. SOM uses the *spherical scoring rule*, which is discussed in the next section.

3.3 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_i, the spherical rule is defined as follows:

$$S\left(\omega_i; \hat{\mu}_i, \hat{\theta}_i\right) = \frac{\mathcal{N}\left(\omega_i; \hat{\mu}_i, \hat{\theta}_i\right)}{\sqrt{\int_{-\infty}^{\infty} \mathcal{N}\left(x; \hat{\mu}_i, \hat{\theta}_i\right)^2 \mathrm{d}x}} \tag{8}$$

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 expected 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.

Using scoring rules to distribute payments to agents in continuous domains is non-trivial. This arises from the fact that $\mathcal{N}(\mu, \theta)$ is unbounded as $\theta \to \infty$. Clearly, if a home agent were to know exactly what it would consume in the next time period, paying it an infinite amount is not acceptable to the aggregator. Therefore, the score must be scaled in order to apply bounds. The next section discusses two methods for achieving this, and introduces the main contribution of this paper – the sum of others' plus max mechanism.

3.4 Sum of Others' plus Max

As has been discussed, scoring rules are used to assign scores to agents based on the accuracy and the precision of the estimates the agents report to the centre. Naturally, agents who report more precise estimates expect to get a higher score, and remembering that the cost agents incur when generating their estimates is proportional to the precision of their generated estimates, rewarding agents based on the score they achieve seems to be a natural development. Considering that we would also like to *fairly* distribute the savings to the agents in a way that is *budget balanced*, one method of distributing the savings to the home agents might be to scale the total savings by the fraction of the sum of all agents' scores that the agent had contributed. We call this mechanism the *percentage contribution* mechanism, where the reward each agent receives is given by:

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

However, despite the fact that this might be intuitively correct, the resulting payments are *not* incentive compatible. Agents are in fact able to *misreport* the precision of their belief in order to gain a higher reward by making their belief seem more precise. In this section, the percentage contribution mechanism is adapted to make the sum of others' plus max (SOM) mechanism, which *is* incentive compatible, but is only *ex ante* weakly budget balanced. That is, in expectation, SOM distributes *at most* 100% of its budget.

This mechanism takes into account not only the spherical score (defined in Equation 8) achieved by the agent, but also those achieved by the other agents in the system. Payments are then determined by multiplying those scores by the savings made by the aggregator when using the reports from the *other* agents in the system, and the aggregator's prior knowledge in place of the report from the agent who is being rewarded. This is necessary in order to preserve incentive compatibility. Furthermore, to ensure agent's payments never outweigh the savings made by the aggregator, it is necessary to provide an upper bound on the precision of reports accepted from agents, 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, the payment agent i obtains, given all agents' estimates is given by:

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

where $\boldsymbol{x}_{-i}=\boldsymbol{x}\setminus\{x_i^*\}$, and the term, $S\left(\omega_i;\omega_i,\theta_{max}\right)$, 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_i when ω_i 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 it's strict propriety, and therefore these payments are incentive compatible, as is shown in Theorem 3.2.

THEOREM 3.2. *SOM is dominant strategy incentive compatible, 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. Ergo the agent is unable to affect the savings used to calculate its payment. Thus, the savings made by the other agents are in effect a constant. Given that, SOM 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. Thus, the expected reward an agent receives is also a unique maximum when it reports truthfully. Therefore, the mechanism is strictly dominant incentive compatible. □

In addition to truth-telling being a strictly dominant strategy, the rewards to agents made by this rule are fairer than those of the uniform mechanism in that the agents are directly compared with each other based upon their score. The reward is simply a scaled fraction of the agents' spherical score over the sum of all other agents' scores. Thus, if an agent scores highly because it reported a precise estimate, and the other agents score lower because of imprecise reports, the first agent will receive a greater share of the savings made. It is essential to divide the agent's spherical score by the sum of the other agents' prescaled scores *plus the maximum score* in order to maintain weak budget balance. Theorem 3.3 provides a proof of the fact that SOM is ex ante weakly budget balanced.

THEOREM 3.3. *SOM is ex ante weakly budget balanced.*

PROOF. Let each agent, i, obtain the score S_i, and $\bar{\Delta}\left(\theta\right)$ be the expected savings made when the agents' reports produce an aggregate precision of θ. In SOM when each agent, i's, report has precision θ_i, the total expected payout is:

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

and the aggregator's total expected savings is

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

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

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

We start with the axiom,

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

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 SOM 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 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 SOM compared to the uniform mechanism is that it is only ex ante weakly budget balanced – a weaker concept than the strict budget balance exhibited by the uniform mechanism. The home agents might 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. This is further explained with the aid of empirical evidence in Section 5.3.

4. THE SOCIAL WELFARE SOLUTION

While the design of the two previous mechanisms assumes that the agents are non-cooperative – that is, they seek only to maximise their own profit – the social welfare solution assumes *cooperation* between the agents. In the social welfare solution, agents ignore their own reward and instead maximise the sum of all agents' utilities within the system. In the case shown in this paper, when maximising social welfare, each agent maximises the following function:

$$U^*(\boldsymbol{x}) = \int \cdots \int_0^\infty \Delta(\boldsymbol{x}, \boldsymbol{x}_a, \boldsymbol{\omega}) - \sum_{\forall i \in N} C(\alpha_i, \theta_i)\ \mathrm{d}\omega_1, \cdots, \mathrm{d}\omega_n \tag{9}$$

The social welfare solution is significant as it provides an upper bound for the social welfare that can be achieved within the system. This result can then be used in order to ascertain the *efficiency* of any mechanism that works in the scenario, where more efficient mechanisms are deemed to be those whose social welfare is closer to that of the maximum social welfare. We use the social welfare result in the next section in order to analyse the efficiency of the sum of others plus max, and uniform mechanisms, and to discover the properties that arise in a cooperative model.

The previous sections have introduced two new mechanisms – the *uniform* mechanism and *sum of others' plus max* – and in doing so have discussed and proved some of the formal properties of said mechanisms. The solution whereby agents act cooperatively to maximise social welfare has also been introduced. The next section uses empirical analysis to further analyse emergent behaviour under equilibrium.

5. EMPIRICAL EVALUATION OF THE MECHANISMS

Given the theoretical properties of the mechanisms discussed in the previous sections, it is clear that agents will truthfully report their information to the aggregator. However, although rational agents will be truthful, they will strategise over the *actual* generated precision of their estimate. Therefore, to further analyse the properties of the mechanisms, we must solve for the equilibrium precision of the agents.

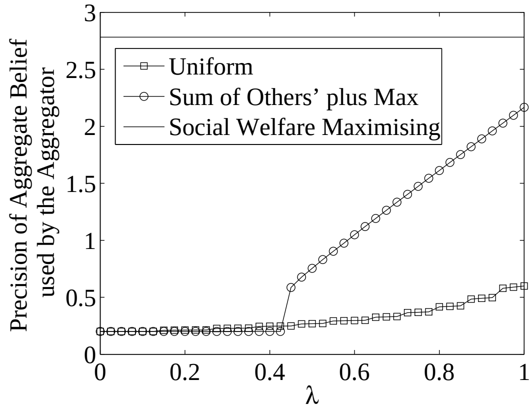

Figure 1: The aggregated precision of the agents' reports, versus the fraction of the savings to be distributed by the aggregator, λ.

Unfortunately, this equilibrium cannot be calculated, and thus, iterated best response is used to computationally find the home agents' equilibrium strategies (the precisions of their generated reports). To aid convergence to an equilibrium, a dynamic named partial best response is used [3]. Equilibrium is detected by measuring the variance of the agents' chosen strategies, and the algorithm is deemed to have reached an equilibrium point when the variance of the last 10 strategies chosen by each agent falls below a predetermined threshold (10^{-8} in our experiments).

Once the equilibria are found, the expected savings and the agents' expected rewards and utilities are analytically calculated. Market values are set to $f = 100$, $\delta^b = 50$, and $\delta^s = 70$. Agents' costs are set to $0.01, 0.02, \ldots, 0.1$, and kept constant throughout each simulation. The aggregator is assumed to have prior information that allows it to make estimates of each houses' demand with precision $\theta_a = 2$.

The remainder of this section discusses the results from these simulations. In Section 5.1, the aggregate precision of all agents' reports are compared for the two mechanisms and the social welfare solution. Next, Section 5.2 discusses the efficiency of each mechanism, where efficiency is defined in terms of a percentage of the social welfare achieved in the cooperative solution. Finally, Section 5.3 discusses the risks to the aggregator in using each mechanism.

5.1 Precision

In this section, the aggregated precisions of the agents' reports under the two mechanisms and the social welfare maximising case are analysed. Figure 1 shows the aggregated precision of all agents' reports against the fraction of the savings the aggregator allocates for distribution. The fact that the agents are playing unique equilibrium strategies that maximise their utility functions means that there is no error in the agents' chosen strategies, and consequently no error in the aggregated precision of the agents' reports. Thus, error bars have been omitted from the plot in Figure 1.

It can be seen from Figure 1 that in terms of aggregated precision, SOM vastly outperforms the uniform mechanism for values of $\lambda > 0.425$, obtaining an aggregated precision which is up to four times the aggregated precision obtained by the uniform mechanism. The step change in the SOM precisions between $\lambda = 0.4$ and $\lambda = 0.5$ occurs because lower values of λ are not incentivising the agents' to produce estimates as their additional reward is outweighed by their costs. Furthermore, steps can be seen in the uniform plot as each agent becomes incentivised to produce a report. The baseline at $\theta = 0.2$ occurs when no agents are incentivised to produce a report with $\theta_i > \theta_{a,i}$. Consequently, the aggregated belief consists only of the aggregator's prior knowledge, thus $\theta = \theta_a/n$ assuming $\theta_a = \theta_{a,i}$, $\forall i \in N$.

Figure 1 also shows the aggregated precision of the reports made by the home agents under a cooperative, social welfare maximising setting. In this setting, each agent's strategy is constant with respect to λ as the agents do not take into account their own reward, as is shown in Equation 9.

Furthermore, note that it is not the absolute value of the reward that is important in determining the strategy to be adopted by the agents, but the gradient of the reward function with respect to the agents' precision, θ_i. The absolute value only becomes relevant if the agent has the choice between the mechanism that is in use. Agents will choose a precision at which the gradient of the reward function, $dP/d\theta_i$ equals the gradient of their cost function $dC_i/d\theta_i$, or α_i. Therefore, it may well be possible to construct a reward function that encourages agents to generate yet more precise estimates than they do with SOM. However, as we discuss in the next section, SOM is already up to 77% efficient, and therefore any increase in the precision generated by agents will only produce marginally improved social welfare.

5.2 Efficiency

The mechanisms discussed in this paper do not 'burn' any unallocated savings. That is, savings that are not distributed to the home agents as payment are not simply discarded. Instead, any unallocated savings are returned to the aggregator to add to its profit. In this way, the mechanisms always allocate 100% of their budget. However, under this model, the budget for each mechanism is itself dependent on the actions of the home agents. Consequently, each mechanism will still result in a different social welfare. With this in mind, in this paper, efficiency is defined as the social welfare achieved by the mechanism expressed as a percentage of the social welfare that is achieved when agents act cooperatively to maximise additional social welfare.

Figure 2 shows the sum of the home agents' additional utilities against the additional utility gained by the aggregator in using each mechanism compared to using no mechanism for values of λ between 0.0 and 1.0. As before, the agents' behaviour in the social welfare solution is independent of λ. It can be seen that SOM is more efficient than the uniform mechanism in that it has points closer to the social welfare maximising solution. Furthermore, for any utility received by the home agents, the aggregator's additional utility is greater under SOM than the uniform mechanism.

Having discussed the expected utilities that are obtained by the agents and aggregator, we see that the aggregator always expects to receive a greater utility when using SOM over the uniform mechanism. However, observing samples of individual rounds shows that there are occasions wherein the aggregator makes a loss. The next section analyses the risk

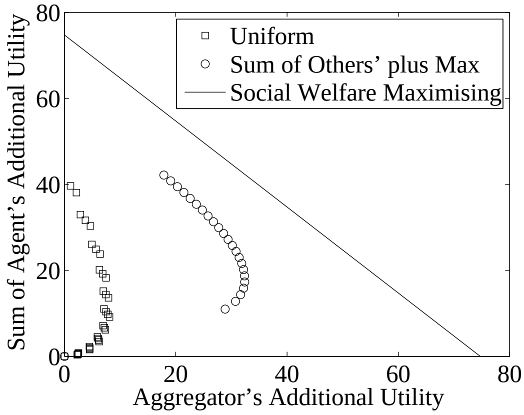

Figure 2: The additional utility of the aggregator when using a mechanism as opposed to using no mechanism against the sum of agents' additional utilities.

to the aggregator in using SOM compared to the benchmark.

5.3 Risk for the Aggregator

This section discusses the risk to the aggregator when using SOM and the uniform mechanism compared to using no mechanism at all – i.e. the aggregator simply using its own beliefs. There is always an element of risk for the aggregator, which arises due to errors made when predicting the demand of the home agents. It should be noted that this work assumes that the agents' beliefs are on average *accurate*. Thus, on average, the savings made in using the agents' more precise information is positive. Nevertheless, it is possible for savings on occasion to be negative, when an agent is confident in its belief, but is incorrect.

However, there are a number of ways the aggregator can mitigate such risk. For one, simply increasing the number of agents being aggregated over decreases the risk to the aggregator, as can be seen by the 'no mechanism' line in Figure 3. The use of mechanisms further reduces the aggregator's risk by encouraging agents to produce more precise estimates.

The results shown in this section use the same experimental setup as above, with the retail price, arbitrarily set to $f_r = f$ (note in a real situation, the aggregator would set $f_r > f$ in order to gain some profit). In addition, the value at risk is calculated through repeated simulation. In more detail, once the agents' strategies have been determined, the simulation of the aggregator purchasing electricity using the agents' reports runs in full for 1000 rounds. In each round, each house, i, is assigned a total consumption, ω_i, sampled from a uniform distribution over the range $[30, 50]$. The agent's own estimate of its consumption is then sampled from a normal distribution around its consumption such that $\mu_i \sim \mathcal{N}(\omega_i, \theta_i)$, where θ_i is the precision chosen by agent i in the equilibrium found through iterated best response as described earlier. The same procedure is used to generate the aggregator's belief using θ_a. The aggregator pays each agent their reward in accordance with the model in Section 2 and mechanisms in Section 3. After each round, the aggregator's utility is calculated (Equation 6). The value of the aggregator's utility at 5% risk is calculated by taking the aggregator's fifth percentile utility from the 1000 rounds. This process is repeated 20 times, and the mean and standard

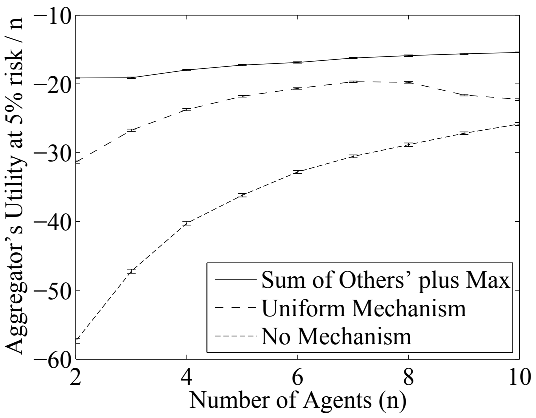

Figure 3: Value of the aggregator's utility per agent at 5% risk for two to ten agents when $\lambda = 0.5$.

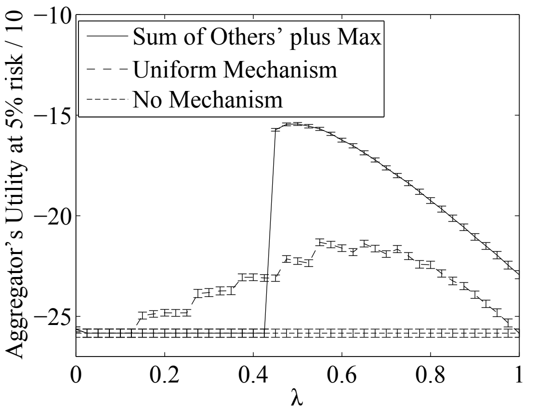

Figure 4: Value of aggregator's utility per agent at 5% risk for ten agents and λ between 0 and 1.

error of the value at 5% risk are taken.

Risk can be further mitigated by the aggregator by adjusting λ, as shown in Figure 4. It can be seen that SOM results in a greatly reduced risk to the aggregator for values of $\lambda > 0.425$ resulting from the high precision estimates that SOM encourages agents to produce. For higher values of λ, although the agents are still producing high precision estimates, as shown in Figure 1, the aggregator is giving away a larger portion of the savings to the agents, thereby decreasing its utility. For values of $\lambda < 0.425$, as shown in Figure 1, agents are not incentivised to produce reports, and thus the risk to the aggregator is equal to the risk when no mechanism is employed. The step changes in the 'uniform' line of Figure 4 arise due to agents individually becoming incentivised to produce estimates for the aggregator.

An additional source of risk comes from the *ex ante* weakly budget balanced nature of SOM; the mechanism is budget balanced *in expectation*, but there may be instances wherein the aggregator makes a loss. Furthermore, in situations whereby the savings made by the aggregator are negative – i.e. the agent's estimates were less accurate than the aggregator's prior estimate – the fact that the mechanism is ex ante, *weakly* budget balanced, means the aggregator will, in expectation, regain *at most* λ of the loss made.

However, it is worth noting that at no point in Figures 3 and 4, is the aggregator worse off by using either mechanism compared to using no mechanism at all. Furthermore, SOM, for $\lambda > 0.425$, maximises the utility of the aggregator. Thus a risk neutral aggregator, who is able to strategise over the mechanism used, will always choose SOM, with $\lambda > 0.425$.

6. CONCLUSIONS

This paper discussed mechanism design in scenarios wherein a centre has some imprecise information regarding a set of values that, when aggregated, provide information it must use to optimally procure goods at a cost. Each variable has a single expert agent that is able, at a cost, to report to the centre more precise information regarding the value. Additionally, the expected cost incurred by the centre increases with the precision of the information it uses to procure said goods. A dominant strategy incentive compatible scoring rule-based mechanism named *sum of others' plus max* (SOM) was developed, which rewards agents from a budget that is equal to the savings made by the centre in using the agents' information over its own.

SOM was compared to a simple mechanism whereby the agents are paid by dividing the savings made equally among the agents. It was shown that the sum of others plus max mechanism increased social welfare compared to the uniform mechanism, and that the social welfare achieved by the sum of others plus max mechanism was 77% that of the optimal solution. Empirical evidence was provided that shows that SOM reduces the risk to the aggregator compared to using a simple uniform mechanism or no mechanism at all.

Future work will investigate the combination of the sum of others plus max mechanism with additional incentives. For example, in the smart grid, demand smoothing is desirable as it reduces the need for extraneous and costly standby generation capacity on the grid. An additional incentive to build into this mechanism would therefore be to reduce the variance of the realised consumptions of each house, ω_i, as agents can potentially gain better payoffs by making their consumptions unpredictable to the aggregator.

7. REFERENCES

[1] G. Brier. Verification of forecasts in terms of probability. *Monthly Weather Review*, 78(1):1—3, 1950.

[2] A. Carvalho and K. Larson. A truth serum for sharing rewards. In *Proc. of AAMAS '11, Taipei, Taiwan*, pages 635—642, 2011.

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

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

[5] R. Nau, V. R. Jose, and R. Winkler. Scoring rules, entropy, and imprecise probabilities. In *Proc. of ISIPTA '07*, pages 307—315, 2007.

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

[7] P. Shi, V. Conitzer, and M. Guo. Prediction mechanisms that do not incentivize undesirable actions. In *Proc. of WINE '09*, pages 89—100, Berlin, Heidelberg, 2009.

[8] R. Winkler. The quantification of judgment: Some methodological suggestions. *Journal of the American Statistical Association*, 62(320):1105—1120, 1967.