

Mechanism Design for Eliciting Probabilistic Estimates from Multiple Suppliers with Unknown Costs and Limited Precision¹

A. Papakonstantinou, A. Rogers, E. H. Gerding and N. R. Jennings

School of Electronics and Computer Science, University of Southampton
{ap06r,acr,eg,nrj}@ecs.soton.ac.uk

Abstract. This paper reports on the design of a novel two-stage mechanism, based on strictly proper scoring rules, that allows a centre to acquire a costly probabilistic estimate of some unknown parameter, by eliciting and fusing estimates from multiple suppliers. Each of these suppliers is capable of producing a probabilistic estimate of any precision, up to a privately known maximum, and by fusing several low precision estimates together the centre is able to obtain a single estimate with a specified minimum precision. Specifically, in the mechanism's first stage M from N agents are pre-selected by eliciting their privately known costs. In the second stage, these M agents are sequentially approached in a random order and their private maximum precision is elicited. A payment rule, based on a strictly proper scoring rule, then incentivises them to make and truthfully report an estimate of this maximum precision, which the centre fuses with others until it achieves its specified precision. We formally prove that the mechanism is incentive compatible regarding the costs, maximum precisions and estimates, and that it is individually rational. We present empirical results showing that our mechanism describes a family of possible ways to perform the pre-selection in the first stage, and formally prove that there is one that dominates all others.

1 INTRODUCTION

The growth of the internet has facilitated large-scale open information systems in which multiple users and providers can autonomously exchange information. Thus, it is important to develop processes that will evaluate this information and will give some guarantees to its quality. This is particularly important in cases where the information in question is a probabilistic estimate or forecast whose generation involves some cost. Examples include online reputation systems where the information in question is the predicted quality of service delivered by a service provider (and the costs are incurred by storing and searching previous interaction histories), and forecasting services where the information is a prediction of some meteorological phenomena (and costs represent the computational expenses of running large scale prediction models). In such cases, there may be incentives for information providers to misreport their estimates or predictions, or to allocate less costly resources to their generation, if they can increase their own utility by doing so (e.g. by being rewarded for more precise estimates than

¹ This research was undertaken as part of the EPSRC funded project on Market-Based Control (GR/T10664/01). This is a collaborative project involving the Universities of Birmingham, Liverpool and Southampton and BAE Systems, BT and HP.

they actually provide). Furthermore, information providers will typically have different capabilities, and thus may only be able to provide estimates up to a certain maximum precision (perhaps due to different computational resources or prediction methodologies). Thus, a potential user of this information is faced with three major challenges. First, it has to elicit the providers' private costs as well as its capabilities, in order to select those which can collectively provide an estimate of the required precision at the minimum cost. In doing so, it may have to fuse together several low precision estimates in order to produce an estimate of sufficiently high precision. Second, it must incentivise these providers to allocate sufficient costly resources into generating this estimate. Finally, it must incentivise the providers to truthfully report their estimates.

Against this background, a number of researchers have proposed the use of *strictly proper scoring rules* to address some of these challenges [1–3]. Mechanisms based on these scoring rules reward precise estimates or forecasts by taking into consideration the difference between an event's predicted and actual outcome (observed at some later stage) when issuing payments. In doing so, they incentivise agents to truthfully report their estimates in order to maximise their expected payment. Miller et al. extend these results and show that by appropriately scaling the scoring rule agents can be incentivised to commit costly resources for generating an estimate of a specific precision [4]. More recently, strictly proper scoring rules have been used to promote the honest exchange of beliefs between agents [5], and within reputation systems to promote truthful reporting of feedback regarding the quality of the experienced service [6].

However, in all the above approaches it is assumed that the costs of the agent providing the estimate or forecast is known by the centre. This is clearly not the case in the examples described above where these costs represent private information known only to each individual agent (since they are dependent on the specific computational resources available to that agent). This issue is addressed by Papakonstantinou et al., who introduce a two-stage mechanism, in which the centre first uses a reverse second price auction to identify a single agent that can provide the estimate with the lower cost, while in the second stage a scaled strictly proper scoring rule incentivises that agent to generate and truthfully report this estimate [7]. While this approach is effective in the specific case it considers, it relies on the assumption that the generation of the estimate is allocated to a single agent, and that this agent can provide an estimate of unlimited precision. In reality, this is unlikely to be the case. Agents will likely have different capabilities and this will be reflected in the maximum precision of the estimate that they can produce. Thus the centre may have to procure estimates from multiple agents and fuse them together to achieve its required precision. This fusion relies on the fact that combining independent unbiased estimates will always result in a fused estimate of greater precision than any of the individual ones (indeed, this relationship is additive in the case of Gaussian estimates [8]).

It is this challenge we address in this paper. Such an extension is challenging since procuring estimates from multiple agents results in an interdependent valuation setting with so-called allocative externalities, for which it has been shown that no standard mechanism exists which is both efficient and incentive compatible [9]. We address this by developing a two stage mechanism whereby the first stage incentivises the agents to truthfully reveal their costs to the centre, by pre-selecting M of the original N agents, and the second stage incentivises a further subset of these M agents to truthfully report

their maximum precisions, and to generate estimates at those precisions. This mechanism is incentive compatible, but not efficient. Thus, we empirically compare several approaches to perform the pre-selection, and identify one that minimises this inefficiency. In more detail, in this paper we extend the state of the art in the following ways:

- We present a two stage mechanism in which a centre in the first stage asks the agents to report their costs and uses these costs to pre-select M agents. In the second stage the centre sequentially asks the pre-selected agents to report their maximum precisions and then to generate an estimate of precision equal to their reported maximum precision. These agents are paid for their estimates using an appropriately scaled strictly proper scoring rule, and this second stage proceeds iteratively until the centre has achieved its required precision (or it has run out of agents, in which case the centre will have to suffice with the closest precision to its required one).
- We prove that our mechanism is incentive compatible in costs, maximum precisions and estimates revealed, and that it is individually rational. Furthermore, we show that the agents maximise their expected utilities by generating estimates of precision equal to their reported maximum precisions.
- We introduce a family of processes by which the centre may pre-select M from N agents. Within this family, the centre divides the agents into groups of $n \leq N$ agents, asks them to reveal their costs, and then selects the m cheapest agents. The $(m+1)^{th}$ cost is then used within the subsequent scoring rule payment. In a setting where costs are linear, we empirically evaluate this family for various values of the parameters n and m , and we calculate the total expected payment made by the centre, and the probability that it actually achieves its required precision, $P(\theta_0)$. We show both empirically and analytically that if the centre forms a single group of agents such that $n = N$ and $m = M$, it minimises its expected total payment.

The rest of this paper is organised as follows: In section 2 we describe our setting in more detail, and in section 3 we present the necessary background on strictly proper scoring rules. In section 4 we detail the mechanism and formally prove its properties, before empirically evaluating the mechanism for various values of its parameters in section 5. We conclude in section 6.

2 ELICITING INFORMATION FROM MULTIPLE SOURCES

We now describe our setting in more detail. We assume that a centre is interested in acquiring a probabilistic estimate or forecast (such as the expected quality of a service within a reputation system or a forecast of the temperature in a weather prediction setting) with a minimum precision θ_0 , henceforth referred to as the required precision. We note that the centre derives no additional benefit if the estimate is of precision greater than θ_0 . Furthermore, we assume that there are $N \geq 2$ rational and risk neutral agents that can provide the centre with an unbiased but noisy estimate or forecast, with mean \hat{x}_i and precision $\hat{\theta}_i$. We model these agents' private estimates as conditionally independent Gaussian random variables such that $x_i \sim \mathcal{N}(x_0, 1/\theta_i)$, where x_0 is the true state of the parameter being estimated¹. The true state, x_0 , is unknown to both the centre and

¹ Note that there is no requirement that the agents should truthfully reveal their true estimates and therefore, we will use a payment rule to incentivise them to do so.

the agents at the time that the estimate is requested, but becomes available to both at some time in the future.

We model the agents' cost of producing their estimate as a function of the precision of the estimate, $c(\theta)$. While the centre has no information regarding the agents' cost functions, we assume that all cost functions are convex (i.e. $c_i''(\theta) \geq 0$), increasing and that $c_i(0) = 0$. We believe that these are realistic assumptions in all cases where there are diminishing returns as the precision increases. We do not assume that all agents use the same cost function, but we do demand that the costs of different agents do not cross or overlap (i.e. the ordering of the costs and their derivatives is the same over all precisions).

Finally, we assume that the maximum precision that each agent can provide is limited and is given by θ_i^c . Thus, agents can produce estimates of any precision up to and including this maximum value (i.e. $0 \leq \theta_i \leq \theta_i^c$). Given this limit, the centre may not be able to rely on a single agent to achieve its required precision, and may have to combine estimates of multiple agents.

To this end, in order to fuse k conditionally independent and unbiased probabilistic estimates, $\{\hat{x}_1, \dots, \hat{x}_k\}$ of possibly different precisions $\{\hat{\theta}_1, \dots, \hat{\theta}_k\}$, into one estimate with mean \bar{x} and precision $\bar{\theta}$, the centre uses the standard result (see [8]):

$$\bar{x}\bar{\theta} = \sum_{i=1}^k \hat{x}_i \hat{\theta}_i \quad \text{and} \quad \bar{\theta} = \sum_{i=1}^k \hat{\theta}_i \quad (1)$$

Note that for this fusion to be appropriate, agents must be incentivised to truthfully report both the means and precisions of their estimates. Now, given this model, the challenge is to design a mechanism in which the centre will be able to identify those agents that can provide their estimates at the lowest cost, and to provide a payment to these agents which will motivate them to truthfully report their maximum precisions and generate and truthfully report their estimates with precisions equal to their reported maximum precisions.

3 STRICTLY PROPER SCORING RULES

As previously discussed, strictly proper scoring rules have been used in computer science to elicit incentive compatible payments and to motivate agents to allocate costly resources into generating a probabilistic estimate or a forecast. However, the existing mechanisms are restricted to cases where either the cost functions are known or the centre is only interested in a single agent's estimate. Before we proceed to the analysis of our mechanism, which relaxes both these assumptions, we provide the essential background regarding strictly proper scoring rules. In brief, such rules are functions that are maximised when agents' predictions of a future event are close to the actual outcomes. Much of the literature concerns three specific rules:

1. Quadratic: $S(x_0|r(x)) = 2r(x_0) - \int_{-\infty}^{\infty} r^2(x)dx$
2. Spherical: $S(x_0|r(x)) = r(x_0)/(\int_{-\infty}^{\infty} r^2(x)dx)^{1/2}$
3. Logarithmic: $S(x_0|r(x)) = \log r(x_0)$

where, $S(x_0|r(x))$ is the payment given to an agent after it has reported its estimate, represented as probability density function $r(x)$, and x_0 is the actual outcome observed.

Payments based upon strictly proper scoring rules not only reward accurate estimates, but also incentivise honest reporting of these estimates. That is, an agent will maximise its expected score, and hence the payment it receives, by reporting its true probabilistic estimate to the centre. Using that result, we can analytically calculate the score an agent expects to receive after generating an estimate of precision θ and truthfully reporting it. The first step is to replace the general density function $r(x)$, with the Gaussian distributions, $\mathcal{N}(x_0; x, 1/\theta)$, and derive the following expressions:

1. Quadratic: $S(x_0; x, \theta) = 2\mathcal{N}(x_0; x, 1/\theta) - \frac{1}{2}\sqrt{\frac{\theta}{\pi}}$
2. Spherical: $S(x_0; x, \theta) = \left(\frac{4\pi}{\theta}\right)^{\frac{1}{4}} \mathcal{N}(x_0; x, 1/\theta)$
3. Logarithmic: $S(x_0; x, \theta) = \log(\mathcal{N}(x_0; x, 1/\theta))$

After integrating over the expected outcome, x_0 , we derive an agent's expected score, $\bar{S}(\theta)$, if it makes, and truthfully reports an estimate of precision θ :

1. Quadratic: $\bar{S}(\theta) = \frac{1}{2}\sqrt{\frac{\theta}{\pi}}$
2. Spherical: $\bar{S}(\theta) = \left(\frac{\theta}{4\pi}\right)^{\frac{1}{4}}$
3. Logarithmic: $\bar{S}(\theta) = \frac{1}{2}\log\left(\frac{\theta}{2\pi}\right) - \frac{1}{2}$

Note that the properties of the strictly proper scoring rules (i.e. the expected score being maximised when agents truthfully report their estimates) are maintained under an affine transformation. Indeed, Miller et al. make use of this observation and use a scaled scoring rule as a payment in order to elicit effort from an agent and to ensure individual rationality [4]. Thus, the agents' expected payment, $\bar{P}(\theta)$, is $\bar{P}(\theta) = \alpha\bar{S}(\theta) + \beta$, where α and β are the scaling parameters.

In more detail, if a centre is interested in acquiring an estimate of precision θ_0 , it can choose the value of α such that the agent's expected utility is maximised when it produces and truthfully reports an estimate of the required precision, θ_0 , such that:

$$\alpha = \frac{c'_t(\theta_0)}{\bar{S}'(\theta_0)} \quad (2)$$

Furthermore, the centre can use the constant β to ensure that it makes the minimum payment to the agent, while still ensuring that the mechanism is individually rational. This is done by ensuring that $\bar{U}(\theta_0) = 0$, thus giving:

$$\beta = c_t(\theta_0) - \frac{c'_t(\theta_0)}{\bar{S}'(\theta_0)}\bar{S}(\theta) \quad (3)$$

However, Miller et al. assumed that the agents' costs c_t are known to the centre. Now, Papakonstantinou et al. relaxed this assumption. They showed that by asking agents to declare their costs in the first stage of a two stage mechanism, selecting the cheapest agent but calculating the scaling parameters using the second lowest revealed cost led to an incentive compatible and individually rational mechanism, in which the

selected agent produces an estimate of precision θ^* , denoted by $\theta^* = \operatorname{argmax}_{\theta} \bar{U}(\theta)$, which is greater or equal to θ_0 (with $\theta^* = \theta_0$ when $c_t = c_s$).

Against this background, we extend the above mechanism so it can deal with the centre's need to combine multiple sources of information in order to get a more accurate estimate. In doing so, we relax the unrealistic assumption of this earlier work, that agents are capable of producing estimates with unlimited precision. In the following section, therefore, we introduce and analyse an iterative two-stage mechanism, which can be applied in the more realistic setting in which agents can only produce estimates up to a privately known maximum precision.

4 THE MECHANISM

In this section, we develop our mechanism so that the precision of the resulting fused estimate is greater than or equal to the centre's required precision. To this end, in the first stage, the centre pre-selects M of the N available agents based on their reported costs. In the second stage, it sequentially asks them to reveal their private maximum precision, in a random order that is independent of their reported costs, until it achieves its required precision, θ_0 . It then incentivises these pre-selected agents to produce and truthfully report their estimates by using an appropriately scaled strictly scoring rule as the basis of each payment.

In more detail, our two-stage iterative mechanism proceeds as follows:

- First Stage
 1. The centre selects $n \geq 2$ agents from the available N and asks them to report their cost functions $\hat{c}_i(\theta)$ with $i \in \{1, \dots, n\}$.
 2. The centre selects the m , ($1 \leq m < n$), agents with the lowest costs, associates the $(m+1)^{th}$ cost with these agents and discards the remaining $n-m$ agents.
 3. The centre repeats the above two steps until it has asked all N agents to report their cost functions. Note that when N is not exactly divisible by n and we have a single remainder, it is discarded. Otherwise in the final round the centre modifies n and m such that $n = N \bmod n$ and $m = \min(m, n-1)$.
 4. We denote the total number of the agents pre-selected in this stage as M and note that its value depends on N , n and m .
- Second Stage
 1. The centre sets its required precision θ_r equal to θ_0 .
 2. The centre randomly selects one of the pre-selected agents and asks it to report its maximum precision $\hat{\theta}_j^c$, with $j \in \{1, \dots, M\}$. It then asks the agent to produce an estimate of this precision and presents this agent with a scaled strictly proper scoring rule. The scaling parameters α and β are determined using equations 2 and 3. However, within these expressions $\hat{\theta}_j^c$ is used instead of θ_0 , and c_s (the cost associated with this agent in the preceding stage – $(m+1)^{th}$ cost in the group from which it was selected) is used instead of c_t . Hence, the scaling parameters are given by:

$$\alpha_j = \frac{c'_s(\hat{\theta}_j^c)}{S'(\hat{\theta}_j^c)} \quad \text{and} \quad \beta_j = c_s(\hat{\theta}_j^c) - \frac{c'_s(\hat{\theta}_j^c)}{S'(\hat{\theta}_j^c)} \bar{S}(\hat{\theta}_j^c) \quad (4)$$

3. The centre sets $\theta_r = \theta_r - \min(\theta_r, \hat{\theta}_j^c)$ and if $\theta_r > 0$ it repeats step two of the second stage.
4. The pre-selected agents produce estimates x_j with precision θ_j and report \hat{x}_j and $\hat{\theta}_j$ to the centre², which after observing the actual outcome, x_0 , issues the following payments to each one of them:

$$P_j(x_0; \hat{x}_j, \hat{\theta}_j) = \alpha_j S_j(x_0; \hat{x}_j, \hat{\theta}_j) + \beta_j \quad (5)$$

with α and β being already determined in step two of the second stage.

In the next section we will prove that this mechanism leads the agents to truthfully reveal their costs in the first stage (so that those which can produce the estimate at the lowest cost can be identified), and that the M pre-selected agents are incentivised to truthfully report their maximum precisions to the centre and subsequently make and truthfully report estimates of these precisions in the second stage. These properties are not obvious, and depend rather subtly on the details of the mechanism. For example, we note that if after asking all M agents for their maximum precisions, the centre does not achieve its required precision, the mechanism must proceed to the payment phase (step 4 in second stage). That is, the centre must commit to paying all pre-selected agents for their estimates at their reported maximum precisions, even if it does not acquire its required precision. Failure to observe this policy would lead agents to over-report their maximum precision, in order that some payment was received, and thus, the mechanism would thus would no longer be incentive compatible in terms of maximum precisions.

Furthermore, note that in step 2 of the first stage, the centre chooses the m agents with the lowest reported costs, and discards the remaining $n - m$ agents. If these agents were not discarded, but were placed back into the pool of available agents, then the mechanism would no longer be incentive compatible in terms of costs; agents would have an incentive to over-report their costs, such that when they are eventually pre-selected, their payment rule will be calculated using a higher cost. Finally, in step 2 of the second stage, the centre must randomly ask the pre-selected agents to report their maximum precisions using an ordering which is independent of their reported costs. Failing to do so will undermine incentive compatibility in terms of costs of the first stage of the mechanism, thereby illustrating how the two stages interact.

5 PROOF OF THE MECHANISMS'S PROPERTIES

Having detailed the mechanism, we now identify and prove its economic properties. Specifically, we show that:

1. The mechanism is incentive compatible with respect to the pre-selected agents' reported maximum precisions and reported estimates.
2. The mechanism is incentive compatible with respect to the agents' reported costs.
3. The mechanism is individually rational.

² Note that we could restrict agents to report their estimates with precision $\hat{\theta}_j^c$. However, as we shall show in section 5, under this mechanism the agents are automatically incentivise to report $\hat{\theta}_j = \hat{\theta}_j^c$ anyway.

Theorem 1 *The mechanism is incentive compatible with respect to the pre-selected agents' reported maximum precisions and reported estimates.*

Proof. Given the mechanism described above, when the agent reports its estimate, it must do so with the precision that it claimed was its maximum. Thus, $\hat{\theta} = \hat{\theta}^c$. Now, given the scaling of the scoring rules described in step 2 in the second stage of the mechanism, the expected utility of the agent, if it reports its maximum precision as $\hat{\theta}^c$, and subsequently produces an estimate of precision θ , which it reports with precision $\hat{\theta}^c$, is denoted by $\bar{U}(\theta, \hat{\theta}^c)$, and given by:

$$\bar{U}(\theta, \hat{\theta}^c) = \frac{c'_s(\hat{\theta}^c)}{\bar{S}'(\hat{\theta}^c)} \left(\bar{S}(\theta, \hat{\theta}^c) - \bar{S}(\hat{\theta}^c) \right) + c_s(\hat{\theta}^c) - c_t(\theta) \quad (6)$$

where $\bar{S}(\theta, \hat{\theta}^c)$ is the agent's expected score for producing an estimate of precision θ and reporting its precision as $\hat{\theta}^c$. Furthermore, $\bar{S}(\hat{\theta}^c)$ is the agent's expected score for producing and truthfully reporting an estimate of precision $\hat{\theta}^c$, $c_t(\cdot)$ is the true cost function of the agent, and $c_s(\cdot)$ is the cost function used to produce the scoring rule (i.e. the $(m+1)$ th lowest revealed cost in the group from which the agent was pre-selected). Taking the first derivative of this expression with respect to $\hat{\theta}^c$ gives:

$$\frac{d\bar{U}(\theta, \hat{\theta}^c)}{d\hat{\theta}^c} = \frac{d}{d\hat{\theta}^c} \left(\frac{c'_s(\hat{\theta}^c)}{\bar{S}'(\hat{\theta}^c)} \right) \left(\bar{S}(\theta, \hat{\theta}^c) - \bar{S}(\hat{\theta}^c) \right) + \frac{c'_s(\hat{\theta}^c)}{\bar{S}'(\hat{\theta}^c)} \bar{S}'(\theta, \hat{\theta}^c) \quad (7)$$

Now, since S is a strictly proper scoring rule, then $\bar{S}(\theta, \hat{\theta}^c) = \bar{S}(\hat{\theta}^c)$ and $\bar{S}'(\theta, \hat{\theta}^c) = 0$ when $\theta = \hat{\theta}^c$. Hence:

$$\left. \frac{d\bar{U}(\theta, \hat{\theta}^c)}{d\hat{\theta}^c} \right|_{\hat{\theta}^c = \theta} = 0 \quad (8)$$

and thus, the utility of the agent is maximised when it reveals as its maximum precision, the precision of the estimate that it subsequently produces³.

We now show that it will actually produce an estimate of precision equal to its maximum precision. To this end, we note that when $\hat{\theta}^c = \theta$, the expected utility of the agent is given by:

$$\bar{U}(\theta) = c_s(\theta) - c_t(\theta) \quad (10)$$

³ For completeness, we confirm that the second derivative is negative at $\theta = \hat{\theta}^c$. To this end, the second derivative is given by:

$$\frac{d^2\bar{U}(\theta, \hat{\theta}^c)}{d(\hat{\theta}^c)^2} (\hat{\theta}^c = \theta) = \frac{c'_s(\hat{\theta}^c)}{\bar{S}'(\hat{\theta}^c)} \bar{S}''(\theta, \hat{\theta}^c) - c''_s(\hat{\theta}^c) + \frac{c'_s(\hat{\theta}^c)}{\bar{S}'(\hat{\theta}^c)} \bar{S}''(\hat{\theta}^c) \quad (9)$$

Now, the first term of equation 9 is negative because S is strictly proper, and this implies that $\bar{S}''(\theta, \hat{\theta}^c)$ is negative at $\theta = \hat{\theta}^c$. Furthermore, $c''_s(\hat{\theta}^c)$ is positive, assuming convexity of the cost function, and $\bar{S}''(\hat{\theta}^c)$ is negative assuming concavity of the scoring rule. Hence, the second derivative is negative at $\hat{\theta}^c = \theta$.

Since $c_s(\cdot)$ and $c_t(\cdot)$ do not cross or overlap, and $c'_s(\theta) > c'_t(\theta)$, then $\bar{U}(\theta)$ is a strictly increasing function. Thus the agent will maximise its expected utility by producing an estimate at its maximum precision, and thus, $\theta = \theta^c$, and hence, $\hat{\theta}^c = \hat{\theta} = \theta^c$, as required. \square

Theorem 2 *The mechanism is incentive compatible with respect to the agents' reported costs.*

Proof. We prove this by contradiction and consider two cases depending on whether or not an agent is pre-selected in the first stage of the mechanism as a result of its misreporting. Let $c_t(\cdot)$ and $\hat{c}(\cdot)$ denote an agents' true and reported cost functions respectively. Furthermore, let $c_s(\cdot)$ denote the cost function used to scale the scoring rule if that agent is among the m agents with the lowest reported costs in its group of n agents in the first stage of the mechanism (i.e. $c_s(\cdot)$ is the $(m+1)^{th}$ cost of that group).

First, suppose that the agent's misreporting does not affect whether it is pre-selected or not. In this case, had the agent been pre-selected, its payment would have been based on the $(m+1)^{th}$ cost of its group and therefore independent of its own report. Conversely, had the agent not been pre-selected, it would have received zero utility, since the remaining $n-m$ agents, of a group of initially n agents, that are not pre-selected are discarded. Hence, there is no incentive to misreport.

Second, suppose that the agent's misreporting affects whether that agent is pre-selected or not. There are now two cases: (1) the agent is pre-selected by misreporting but would have not been if it was truthful, (i.e. $c_t(\hat{\theta}^c) > c_s(\hat{\theta}^c)$ and $\hat{c}(\hat{\theta}^c) < c_s(\hat{\theta}^c)$), and (2) the agent is not pre-selected by misreporting but would have been if truthful (i.e. $c_t(\hat{\theta}^c) < c_s(\hat{\theta}^c)$ and $\hat{c}(\hat{\theta}^c) > c_s(\hat{\theta}^c)$).

Case (1). Since the true cost $c_t(\hat{\theta}^c) > c_s(\hat{\theta}^c)$, it follows directly from theorem 1 that the expected utility $\bar{U}(\theta) = c_s(\theta) - c_t(\theta)$ is strictly negative, irrespective of θ . Therefore, the agent could do strictly better by reporting truthfully in which case the expected utility is zero.

Case (2). In this case the agent would have been pre-selected if it was truthful, but now receives a utility of zero since it hasn't been pre-selected due to its misreporting. To show that this type of misreporting is suboptimal, we need to show that, when $c_t(\hat{\theta}^c) < c_s(\hat{\theta}^c)$, an agent benefits from being pre-selected, since it may then be asked to generate an estimate at its reported maximum precision, $\hat{\theta}^c$. It follows directly from theorem 1 that $\bar{U}(\hat{\theta}^c) = c_s(\hat{\theta}^c) - c_t(\hat{\theta}^c) > 0$ when $c_t(\hat{\theta}^c) < c_s(\hat{\theta}^c)$, and therefore there is no incentive for an agent that would have been pre-selected to misreport its cost function. \square

Theorem 3 *The mechanism is interim individually rational.*

Proof. Due to theorem 2, we can assume that all agents, and consequently those pre-selected, will report their true cost functions, and therefore $c_t(\theta) \leq c_s(\theta)$. In theorem 1 we show that the expected utility $\bar{U}(\theta) = c_s(\theta) - c_t(\theta)$ is strictly non-negative, irrespective of θ . Therefore, the expected utility of a pre-selected agent that generates an estimate of precision equal to its reported maximum precision $\hat{\theta}^c$, is strictly non-negative (i.e. $\bar{U}(\hat{\theta}^c) \geq 0$), and hence the mechanism is *interim* individually rational. Note that

the mechanism is *interim* individually rational, since the utility is non-negative in expectation but there may be instances, in which the payment could be negative if the prediction turns out to be far from the actual outcome. Setting β in the second stage, such that the payments are always positive, would make the mechanism ex-post individually rational. However, this would then violate the incentive-compatibility property since the agents could receive positive payoffs by misreporting their cost functions in the first stage. \square

6 EMPIRICAL EVALUATION

Having proved the economic properties of the mechanism, we present empirical results for a specific scenario in order to explore the effect that the parameters n and m have on the centre’s total payments, and whether or not it obtains its required precision. In more detail, in this scenario the cost functions are represented by linear functions, given by $c_i(\theta) = c_i\theta$, where c_i are independently drawn from a uniform distribution $c_i \sim \mathcal{U}(1, 2)$. The maximum precisions of the selected agents, θ_i^c , are independently drawn from another uniform distribution $\theta_i^c \sim \mathcal{U}(0, 1)$ and finally the centre’s required precision, θ_0 , is equal to 1.7 in order to generate representative results whereby the probability of achieving the required precision, $P(\theta_0)$, covers a broad range of values in $[0, 1]$. Finally, we restrict our analysis to the use of spherical scoring rule, as Papanikolaou et al. have shown that among the quadratic, spherical and logarithmic rules, this scoring rule is the most appropriate choice for a centre that is attempting to minimise its payment [7].

To this end, for $N = 7$ we explore all possible combinations of n and m given the constraints that $2 \leq n < N$ and $1 \leq m < n$. For each combination, we simulate the mechanisms 10^7 times and for each iteration we record whether the centre was successful in acquiring an estimate at its required precision, and the sum of all the payments it issued to those agents who were asked to produce an estimate. In figure 1 we plot, for each possible combination of n and m , the probability of acquiring the required precision and the total payment made by the centre. We note that given the number of iterations that we perform, the standard error in the mean values plotted are much smaller than the symbol size shown in the plot, and thus, we omit them.

We show the results of these experiments in figure 1, in which each possible combination of n and m is indicated with an identical marker. However, we differentiate the case where $n = N$ and the reason for this shall become clear shortly. Furthermore, we also indicate the case where the centre has full information of the agents’ costs and maximum precisions, and may thus optimally select the agents it needs in order to ensure that θ_0 is achieved at the minimum cost. This case represents an upper bound for the mechanism, since the centre can now maximise its probability of achieving its required precision by selecting all the available N agents.

Now, regarding figure 1, we first note that many possible combinations of n and m give rise to the same value of $P(\theta_0)$, and thus the family of possible pre-selection methods fall into 6 distinct columns. This is because this probability depends only on the number of agents that are pre-selected (denoted by M) and many of these combinations result in the same number of agents being pre-selected (i.e. if $N = 7$, both $n = 4$,

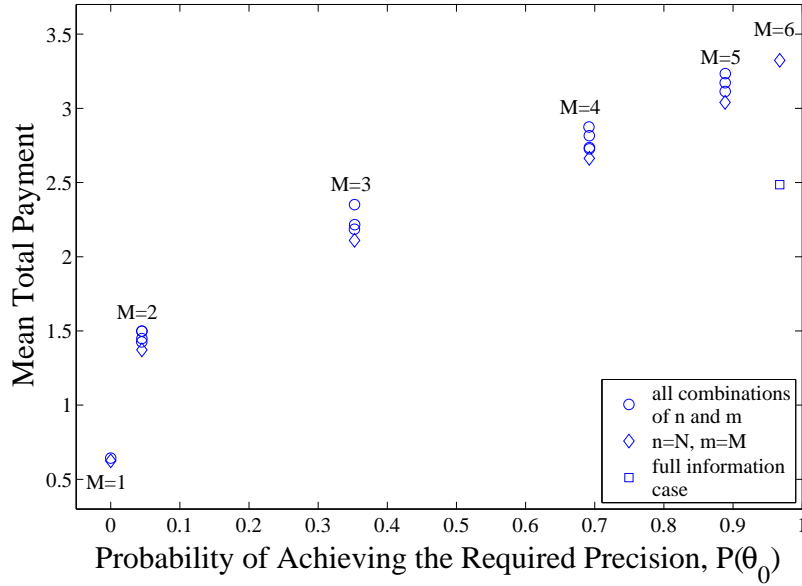


Fig. 1. Centre's probability of achieving the required precision and the mean total payment it has to issue.

$m = 2$ and $n = 5$, $m = 3$ result in $M = 4$). In the following theorem, we show that this probability increases, as the number of the pre-selected agents increases:

Theorem 4 *The probability the centre has of achieving its required precision increases as M increases.*

Proof. We consider the sum Θ of M independent and uniformly distributed random variables $\theta_i^c \sim \mathcal{U}(0, 1)$ which denote the agents' maximum precisions. If Θ is a random distribution such that $\Theta = \theta_1^c + \dots + \theta_M^c$, its cumulative distribution allows us to calculate $P(\Theta \geq \theta_0)$ as follows:

$$P(\Theta \geq \theta_0) = \begin{cases} 1 - \frac{1}{M!} \sum_{i=0}^{\lfloor \theta_0 \rfloor} (-1)^i \binom{M}{i} (\theta_0 - i)^M & 0 \leq \theta_0 \leq M \\ 0 & \theta_0 > M \end{cases} \quad (11)$$

and it is easy to verify that this is an increasing function in M . \square

Second, note that for each possible value of M , the case where $n = N$ and $m = M$ dominates all other combinations of n and m (i.e. it results in the lowest mean total payment). This case corresponds to a single selection stage in which M agents are pre-selected directly from the original N in a single step. This result is demonstrated empirically in figure 1, and we present an analytical proof for the specific case that we consider here:

Theorem 5 *In a setting with linear cost functions, where agents' costs and maximum precisions are independently drawn from uniform distributions, for a given probability of achieving θ_0 , the centre minimises its expected total payment when $n = N$ and $m = M$.*

Proof. Given the mechanism and setting described above, we first note that when the costs of the agents are represented by linear functions, then $c_i(\theta) = c_i\theta$, and hence, $c'_i(\theta) = c_i$. Using this result within the scaling parameters of the payment rule described in step two of the second stage of the mechanism, gives the result:

$$\alpha_j = \frac{c_s}{\overline{S}'(\widehat{\theta}_j^c)} \quad \text{and} \quad \beta_j = c_s \widehat{\theta}_j^c - \frac{c_s}{\overline{S}'(\widehat{\theta}_j^c)} \overline{S}(\widehat{\theta}_j^c) \quad (12)$$

Thus, both α and β are proportional to c_s , and hence the payment to any agent is also proportional to the cost used in the calculation of the scaling parameters. Secondly, we note that due to the random selection of agents within the second stage of the mechanism, the precision of the estimate generated by any agent is independent of the cost used to generate its payment rule. Hence, the expected total payment to the agents is proportional to the mean cost used to generate their payment rules.

Now, the costs used to generate the payment rule of any agent is the $(m + 1)^{th}$ lowest reported cost when m agents are pre-selected from n . Thus, in order to show that setting $n = N$ and $m = M$ minimises the expected total payment of the centre, we must simply show that the expected value of the $(M + 1)^{th}$ cost when pre-selecting M agents from N , is lower than any other combination. To do so, we note that if the costs of the agents are i.i.d. from the standard uniform distribution⁴, and the agents report truthfully (as they are incentivised to do), then the density function that describes the $(m + 1)^{th}$ cost, denoted by $C_{m+1:n}$, is given by:

$$c_{m+1:n}(u) = \frac{n!}{m!(n-m-1)!} u^m (1-u)^{n-m-1}, \quad 0 \leq u \leq 1 \quad (13)$$

and Arnold et al. show that the mean of this distribution is simply $\frac{m+1}{n+1}$ [10]. Thus, we now prove that the $(M + 1)^{th}$ cost when pre-selecting all M agents directly from N is less than the expected cost that results from first pre-selecting m agents from n and then pre-selecting the remaining $M - m$ agents from $N - n$. Thus, we must prove the inequality:

$$\left(\frac{M+1}{N+1} \right) \leq \frac{m}{M} \left(\frac{m+1}{n+1} \right) + \frac{M-m}{M} \left(\frac{M-m+1}{N-n+1} \right) \quad (14)$$

subject to the constraints that $M < N$, $m < n$ and $N - n > M - m$, and we note that if it holds in this case, then it holds for all possible combinations of n and m .

A first step towards the proof of equation 14, is performing the following substitutions:

$$a = m, \quad b = M - m, \quad c = n, \quad d = N - n$$

⁴ For notational simplicity we shall assume that the costs are drawn from $\mathcal{U}(0, 1)$, but we note that the proof is valid for a uniform distribution of any support.

and now equation 14 takes the following form:

$$\frac{(a+b)(a+b+1)}{c+d+1} \leq \frac{a(a+1)}{c+1} + \frac{b(b+1)}{d+1} \quad (15)$$

with $a, b, c, d \geq 0$, $a < c$, and $b < d$.

Now, by multiplying all fractions in equation 15 to obtain a common denominator, $(c+1)(d+1)(c+d+1)$, and noting that this denominator is positive, translates equation 15 into the following condition:

$$\begin{aligned} & (a+b)(a+b+1)(c+1)(d+1) \\ & - a(a+1)(c+d+1)(d+1) \\ & - b(b+1)(c+d+1)(c+1) \leq 0 \end{aligned} \quad (16)$$

We can rearrange this expression into the form:

$$F_1(a, b, c, d) + F_2(a, b, c, d) + F_3(a, b, c, d) \leq 0 \quad (17)$$

where:

$$F_1(a, b, c, d) = -(d \cdot a - b \cdot c)^2 \quad (18)$$

$$F_2(a, b, c, d) = -b(c-a)^2 - b^2(c-a) - a(d-b)^2 - a^2(d-b) \quad (19)$$

$$F_3(a, b, c, d) = a(b-d) + b(a-c) \quad (20)$$

Now, it is easy to verify that F_1 , F_2 , and F_3 are all negative given the initial constraints that $a, b, c, d \geq 0$, $a < c$ and $b < d$. Hence, it follows that equation 17 is negative. \square

Finally, we note that the full information case (shown in figure 1 as a square) results in the highest probability of achieving the required precision. This is expected as in this case the centre has the option selecting all N agents to generate estimates, whereas the requirement to pre-select agents to incentivise them to truthfully report their costs limits our mechanism to using at most $N - 1$ agents. We also note that the full information case results in significantly lower total payments to the agents since in this case the centre is able to select those agents with the lowest costs to generate the estimates, and it need only pay them sufficient to ensure that their expected utility is infinitesimally greater than zero. Given that it is always preferable to set $n = N$, the choice of the value of m is determined by the trade-off between the total payment made by the centre and the probability of it acquiring an estimate of its required precision. If the distributions of cost and maximum precisions are known, this can be evaluated prior to running the mechanism through simulation. However, if these distributions are unknown, setting $m = N - 1$ ensures that the probability of acquiring the required precision is maximised.

7 CONCLUSIONS

In this paper we introduced a two-stage mechanism based on strictly proper scoring rules. In this mechanism a centre pre-selects M from the N available agents by eliciting

their cost functions in the first stage. Then, in the second stage, it approaches these M agents and asks them to report their maximum precision and make a costly probabilistic estimate or forecast of that precision. We applied this mechanism in a setting where an agent might not be able to generate estimates of a sufficient precision to individually meet the centre's needs, hence leaving the centre no option but to combine multiple such estimates and fuse them into a more accurate one. Furthermore, in the setting we considered, the agents, and initially the centre, have no information about the cost functions or the maximum precisions. For this setting, we proved that this mechanism is incentive compatible and individually rational. Furthermore, we empirically evaluated the mechanism for various values of the parameters m and n and showed that for a given probability, $P(\theta_0)$, the centre minimises its mean total payment if it pre-selects M agents directly from a single group of N agents.

In our future work we intend to relax the assumption that the centre can observe the actual outcome after receiving the estimates of the selected agents. In doing so, we will be able to address a larger set of problems where observing the actual outcome is impossible or in settings in which the payment must be made before the actual outcome can be observed. In this case, we postulate that we can score each agent against the fused estimates of all the other agents, and thus the dominant truthful reporting strategies of the agents within the mechanism that we presented here are likely to be replaced by a Nash equilibrium in which reporting the truth is optimal if others are also reporting the truth. In this case, we expect the total payment made to the agents to increase, and this loss by the centre reflects the uncertainty in the true outcome, against which the agents must be scored.

References

1. Hendrickson, A.D., Buehler, R.J.: Proper scores for probability forecasters. *The Annals of Mathematical Statistics* **42**(6) (1971) pp. 1916–1921.
2. Savage, L.J.: Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association* **66**(336) (1977) pp. 783–801.
3. Selten, R.: Axiomatic characterization of the quadratic scoring rule. *Experimental Economics* **1**(1) (1998) pp. 43–61.
4. Miller, N., Resnick, P., Zeckhauser, R.: Eliciting honest feedback: The peer prediction method. *Management Science* **51**(9) (2005) pp. 1359–1373.
5. Zohar, A., Rosenschein, J.S.: Mechanisms for information elicitation. *Artificial Intelligence* **172**(16-17) (2008) pp. 1917–1939.
6. Jurca, R., Faltings, B.: Reputation-based service level agreements for web services. In: *In Service Oriented Computing*, Amsterdam (2005) pp. 396–409.
7. Papakonstantinou, A., Rogers, A., Gerding, E.H., Jennings, N.R.: A truthful two-stage mechanism for eliciting probabilistic estimates with unknown costs. In: *Proceedings of the 18th European Conference on Artificial Intelligence*, Patras, Greece (2008) pp. 448–452.
8. DeGroot, M.H., Schervish, M.J.: *Probability and Statistics*. Addison Wesley (2002).
9. Jehiel, P., Moldovanu, B.: Efficient design with interdependent valuations. *Econometrica*, **69**(5) (2001) pp. 1237-1259.
10. Arnold, B.C., Balakrishnan, N., Nagaraja, H.N.: *A First Course in Order Statistics*. SIAM (2008).