

A Truthful Two-Stage Mechanism for Eliciting Probabilistic Estimates with Unknown Costs

Athanasios Papakonstantinou and Alex Rogers and Enrico H. Gerding and Nicholas R. Jennings¹

Abstract. This paper reports on the design of a novel two-stage mechanism, based on strictly proper scoring rules, that motivates selfish rational agents to make a costly probabilistic estimate or forecast of a specified precision and report it truthfully to a centre. Our mechanism is applied in a setting where the centre is faced with multiple agents, and has no knowledge about their costs. Thus, in the first stage of the mechanism, the centre uses a reverse second price auction to allocate the estimation task to the agent who reveals the lowest cost. While, in the second stage, the centre issues a payment based on a strictly proper scoring rule. When taken together, the two stages motivate agents to reveal their true costs, and then to truthfully reveal their estimate. We prove that this mechanism is incentive compatible and individually rational, and then present empirical results comparing the performance of the well known quadratic, spherical and logarithmic scoring rules. We show that the quadratic and the logarithmic rules result in the centre making the highest and the lowest expected payment to agents respectively. At the same time, however, the payments of the latter rule are unbounded, and thus the spherical rule proves to be the best candidate in this setting.

1 INTRODUCTION

In a world where information can be distributed over systems owned by different stakeholders and accessed by multiple users, 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 estimates of quality of service within a reputation system, or forecasts of future events such as weather conditions, where such costs could represent the computational task of accessing and evaluating previous interactions records, or that of running a large scale weather prediction model. Now, when the provider of such information is a rational selfish agent, it may have an incentive to misreport its estimate, or to allocate less costly resources to its generation, if it can increase its own utility by doing so (e.g. by being rewarded for a more precise estimate than it actually provides). Thus, a centre attempting to elicit such information is presented with three challenges. First, it must identify the agent who can provide an estimate of the required precision at the lowest cost. Second, it must incentivise this agent to allocate sufficient costly resources in order to provide an estimate of the required precision. Finally, it must incentivise this agent to truthfully report the estimate that has been generated.

Against this background, a number of researchers have proposed the use of ‘*strictly proper scoring rules*’ to address these challenges

[1, 5]. Mechanisms using these rules reward accurate estimates or forecasts by making a payment to agents based on the difference between an event’s predicted and actual outcome (observed at some later stage). Such mechanisms have been shown to incentivise agents to truthfully report their estimates in order to maximise their expected payment [6]. More recently, strictly proper scoring rules have been used in computer science to promote the honest exchange of beliefs between agents [7], and within reputation systems to promote truthful reporting of feedback regarding the quality of a service experienced [2]. Furthermore, Miller et al. have shown that when the agents’ costs are known, it is possible to use an appropriately scaled strictly proper scoring rule to induce agents to commit costly resources to generate estimates of any required precision [4].

While these approaches are effective in the specific cases that they consider, they all rely on the fact that the cost of the agent providing the estimate or forecast is known by the centre. This is not the case in our scenario where these costs represent private information known only to each individual agent (since they are dependent on the specific computational resources available to the agent). Thus, in addressing this shortcoming, we contribute to the state of the art by presenting a novel two-stage mechanism which relaxes this assumption. The first stage of the mechanism incentivises agents’ to truthfully reveal their costs to the centre, thus allowing it to select the agent with the lowest cost. The second stage then incentivises this agent to generate an estimate with a minimum required precision, and to truthfully report this estimate to the centre. In more detail, in this paper we extend the state of the art in the following ways:

- We describe a novel two-stage mechanism in which a centre uses a reverse second price auction in the first stage to elicit the true costs of agents, and hence identify the agent that can provide an estimate with a specified precision at the lowest cost. An appropriately scaled strictly proper scoring rule is then used in the second stage of the mechanism to incentivise this agent to generate and truthfully report the estimate.
- We formally prove that this mechanism is incentive compatible in both costs and estimates revealed, and that it is individually rational. That is, agents will truthfully report both costs and estimates to the centre, and willingly participate within the mechanism.
- We empirically evaluate our mechanism by comparing the quadratic, spherical and logarithmic scoring rules in a setting where costs depend linearly on precision. We show that while the logarithmic rule results in the centre making the lowest expected payment to the agent, this payment is unbounded. The other rules are bounded, but result in higher expected payments. Hence, we find that the spherical rule is preferred in our setting.

The rest of this paper is organised as follows: In section 2 we describe our model, and in section 3 we present background on strictly

¹ School of Electronics and Computer Science, University of Southampton, Southampton, SO17 1BJ, UK, email: {ap06r,acr,eg,nrj}@ecs.soton.ac.uk

Table 1. Comparison of Quadratic Spherical and Logarithmic Scoring Rules

Scoring Rule:	Quadratic	Spherical	Logarithmic
$S(x_0; x, \theta)$	$2\mathcal{N}(x_0; x, 1/\theta) - \frac{1}{2}\sqrt{\frac{\theta}{\pi}}$	$(\frac{4\pi}{\theta})^{\frac{1}{4}} \mathcal{N}(x_0; x, 1/\theta)$	$\log(\mathcal{N}(x_0; x, 1/\theta))$
$\bar{S}(\theta)$	$\frac{1}{2}\sqrt{\frac{\theta}{\pi}}$	$(\frac{\theta}{4\pi})^{\frac{1}{4}}$	$\frac{1}{2}\log(\frac{\theta}{2\pi}) - \frac{1}{2}$
$\bar{S}'(\theta)$	$\frac{1}{4\sqrt{\pi\theta}}$	$\frac{1}{4}(\frac{1}{4\pi\theta^3})^{\frac{1}{4}}$	$\frac{1}{2\theta}$
α	$4c'(\theta_0)\sqrt{\pi\theta_0}$	$4c'(\theta_0)(4\pi\theta_0^3)^{\frac{1}{4}}$	$2c'(\theta_0)\theta_0$
β	$c(\theta_0) - 2\theta_0c'(\theta_0)$	$c(\theta_0) - 4\theta_0c'(\theta_0)$	$c(\theta_0) - 2c'(\theta_0)\theta_0\left(\frac{1}{2}\log(\frac{\theta_0}{2\pi}) - \frac{1}{2}\right)$

proper scoring rules. In section 4 we detail our mechanism and formally prove its economic properties, before empirically evaluating it in section 5. We conclude and discuss future work in section 6.

2 INFORMATION ELICITATION PROBLEM

We now describe our model in more detail. Specifically, we assume that there is a centre interested in acquiring a probabilistic estimate or forecast (such as an expected quality of service within a reputation system, or a forecast temperature in a weather prediction setting) with a minimum precision θ_0 , henceforth referred to as the required precision². We assume that there are $N \geq 2$ rational, risk neutral agents who can provide the centre with an unbiased but noisy estimate or forecast, x , of precision θ . We model the agents' private estimates as Gaussian random variables such that $x \sim N(x_0, 1/\theta)$, where x_0 is the true state of the parameter being estimated. Note that this true state is unknown to both the centre and the agents at the time that the estimate is requested, but becomes available to the centre at some time in the future. For example, in a reputation system the actual quality of service received is only known once the service has been procured, and in a weather forecasting setting the actual weather that occurs is observed by the centre at some later date.

The agents incur a cost in producing their estimate, and we assume that this cost is 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. $c''(\theta) \geq 0$), and we note that this is a realistic assumption 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 (i.e. the cost ordering of agents is the same over all precisions).

Given this model, the challenge is to design a mechanism that enables the centre to identify the agent that can provide the estimate or forecast at the lowest cost, and to provide a payment to this agent such that it is incentivised to generate the estimate or forecast with a precision at least equal to the required one and to report it truthfully.

3 STRICTLY PROPER SCORING RULES

As discussed in the introduction, the problem described above has previously been addressed through the use of strictly proper scoring rules as payments in the case that the agents' cost functions are known to the centre [2, 4]. Before we proceed to the analysis of our mechanism which is designed for cases where the centre has no knowledge about the costs, we give a brief description of strictly proper scoring rules. As described earlier, such rules are used to calculate a payment to an agent depending on the difference

² Note that we assume that the centre derives no additional benefit if the estimate is of precision greater than θ_0 .

between an event's predicted and actual outcome. Much of the literature of strictly proper scoring rules concerns three specific rules, the quadratic, spherical and logarithmic rules, given by:

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

In each case, $S(x_0|r(x_0))$ 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.

3.1 An Incentive Compatible Mechanism

It is a standard property of strictly proper scoring rules that an agent will maximise its expected score (and hence the payment it receives) by reporting its true probabilistic estimate to the centre [1, 3]. Thus, mechanisms based upon them are incentive compatible. Using this result, we can calculate the score that the agent expects to receive, given that it has generated an estimate of precision θ and has truthfully reported it to the centre (as it is incentivised to do).

To do so, we first note that, in our case, where estimates are represented by Gaussian distributions, we can replace $r(x_0)$ with $\mathcal{N}(x_0; x, 1/\theta)$, and derive new expressions for each of the three scoring rules shown above (these are presented in the first row of table 1). We can then simply integrate over the expected outcome to derive the agents expected score, $\bar{S}(\theta)$. These results are shown in the second row of table 1, and form the basis of the calculations and proofs that we present in the following sections.

3.2 Eliciting Effort with Known Costs

It should now be noted that the above scoring rules will still be incentive compatible if they undergo an affine transformation. Indeed, Miller et al. show that by using appropriate scaling parameters, and given knowledge of an agent's costs, it is possible to induce an agent to make and truthfully report an estimate with a specified precision, θ_0 [4]. In this case, an agent's expected payment, $\bar{P}(\theta)$, is given by:

$$\bar{P}(\theta) = \alpha\bar{S}(\theta) + \beta \quad (1)$$

and the expected utility of the agent is given by:

$$\bar{U}(\theta) = \alpha\bar{S}(\theta) + \beta - c(\theta) \quad (2)$$

The centre can now choose the value of α such that the agent's utility (its payment minus its costs) is maximised when it produces and truthfully reports an estimate of the required precision, θ_0 . To do so, it solves $d\bar{U}/d\theta|_{\theta_0} = 0$ to give:

$$\alpha = \frac{c'(\theta_0)}{\bar{S}'(\theta_0)} \quad (3)$$

In rows three and four of table 1 we present this result, and the derivative of the expected score that is required to calculate it, for each of the three strictly proper scoring rules presented earlier.

3.3 An Individually Rational Mechanism

Finally, we now note that in order for an agent to incur the cost of producing an estimate, it must expect to derive positive utility from doing so. Thus, 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. When costs are known, the centre can do so by making the agents indifferent between producing the estimate or not, by ensuring that $\bar{U}(\theta_0) = 0$, thus giving:

$$\beta = c(\theta_0) - \frac{c'(\theta_0)}{\bar{S}'(\theta_0)} \bar{S}(\theta) \quad (4)$$

Again, row five of table 1 shows this result for each scoring rule.

4 TRUTH ELICITATION MECHANISM FOR UNKNOWN COSTS

In the previous section we discussed how the centre can motivate agents to make a probabilistic estimate or a measurement of a specific precision. However, this analysis assumed the agents' costs are known. In this section we relax this assumption and present a novel two-stage mechanism which first incentivises the agents to reveal their true costs to the centre, and then, based on this information, induces an agent to produce an estimate of *at least* the required precision. In more detail, in the first stage the centre asks the agents to submit their cost functions and then it assigns the estimation task to the agent with the lowest cost. Then, in the second stage, the centre uses a strictly proper scoring rule as before, but now uses the second-lowest cost reported by the agents to scale the scoring rule (i.e., set α and β). This is akin to a reverse second-price or Vickrey auction, where the agents' rewards are equal to the second-lowest reported costs. However, in this case the reward is determined by the scoring rule, and hence depends on the actual estimate produced. In particular, this requires the scaling parameters α and β to be chosen carefully in order to incentivise the agents to reveal their true costs in the first stage. In more detail, our mechanism proceeds as follows:

1. First Stage

- The centre announces that it needs an estimate of required precision θ_0 , and asks all agents $i \in \{1, \dots, N\}$, where $N \geq 2$, to report their cost functions $\hat{c}_i(\theta)$.³
- The centre assigns the forecast or estimate to the agent who reported the lowest cost at the required precision, i.e., agent i such that $\hat{c}_i(\theta_0) = \min_{k \in \{1, \dots, N\}} \hat{c}_k(\theta_0)$.

2. Second Stage

- The centre announces a scoring rule $\alpha S(x_0; x, \theta) + \beta$, where: (1) $S(x_0; x, \theta)$ is a strictly proper scoring rule, (2) $\bar{S}(\theta)$ is strictly concave as a function of precision θ ,⁴ and (3) α and β are determined using equations 3 and 4 respectively, but now based on the second-lowest reported cost functions (i.e. $\hat{c}_j(\theta)$ such that $\hat{c}_j(\theta_0) = \min_{k \neq i} \hat{c}_k(\theta_0)$).

³ We note that in practise the centre only requires $\hat{c}_i(\theta_0)$ and $c'_i(\theta_0)$. However, for notational convenience we request the agents to reveal their entire cost function.

⁴ We note that the quadratic, spherical, and logarithmic scoring rules satisfy both of these properties (see row 2 of table 1).

- The agent selected in the first stage produces an estimate x with precision θ and reports \hat{x} and $\hat{\theta}$ to the centre.
- Once the actual outcome has been observed, the centre then gives the following payment to the agent:

$$P(x_0; \hat{x}, \hat{\theta}) = \alpha S(x_0; \hat{x}, \hat{\theta}) + \beta \quad (5)$$

4.1 Economic Properties of the Mechanism

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

1. The mechanism is incentive compatible in the first stage w.r.t. the costs. Specifically, truthful revelation of agents' cost functions is a weakly dominant strategy.
2. The mechanism is incentive compatible w.r.t. the selected agent's reported measurement and precision in the second stage.
3. The mechanism is individually rational.
4. The centre motivates the selected agent to make an estimate with a precision which is at least as high as θ_0 , the precision required by the centre. We refer to actual precision produced as the 'optimal precision' (from the perspective of the agent) θ^* .

We now formally prove these properties. To do so, we first derive two lemmas which are then used in the proofs that follow. The first lemma shows that, if the true costs of the agent performing the measurement are less than the costs which are used to scale the scoring rule, the optimal precision θ^* will be greater than θ_0 . Let these cost functions be denoted by $c_t(\theta)$ and $c_s(\theta)$ respectively. More formally:

Lemma 1. *If $c_t(\theta_0) < c_s(\theta_0)$, where $c_t(\theta)$ is the agent's true cost function, and $c_s(\theta)$ is the cost function used to scale the scoring function, then $\theta^* > \theta_0$.*

Proof. By scaling the scoring function using equations 3 and 4 and $c_s(\theta)$, the agent's expected utility becomes:

$$\bar{U}(\theta) = \frac{c'_s(\theta_0)}{\bar{S}'(\theta_0)} (\bar{S}(\theta) - \bar{S}(\theta_0)) + (c_s(\theta_0) - c_t(\theta)) \quad (6)$$

Now, the optimal precision θ^* which maximises his expected utility is formally denoted by $\theta^* = \operatorname{argmax}_{\theta} \bar{U}(\theta)$. Therefore, $\bar{U}'(\theta^*) = 0$, and thus we have:

$$\frac{\bar{S}'(\theta^*)}{\bar{S}'(\theta_0)} = \frac{c'_t(\theta^*)}{c'_s(\theta_0)}. \quad (7)$$

Let $f(\theta) = \bar{S}'(\theta)/\bar{S}'(\theta_0)$ and $g(\theta) = c'_t(\theta)/c'_s(\theta_0)$. Since $\bar{S}(\theta)$ is (strictly) concave it is easy to show that $f'(\theta) \leq 0$ for $\theta \geq \theta_0$ and $f'(\theta) < 0$ for $\theta > \theta_0$. Furthermore, since $c_t(\theta)$ is convex $g'(\theta) \geq 0$ for $\theta \geq \theta_0$. Now, since f is decreasing and g is increasing, when $c_t(\theta_0) = c_s(\theta_0)$ clearly the only point which satisfies equation 7 is where $\theta^* = \theta_0$. If $c_t(\theta_0) < c_s(\theta_0)$, on the other hand, it is easy to verify that $g(\theta_0) < 1$, since we assumed the cost functions to be non-crossing. Hence, since $f(\theta_0) = 1$, the only solution where the two function meet is where $\theta > \theta_0$, and thus, $\theta^* > \theta_0$. \square

The next lemma shows that, if the true costs of the agent doing the measurement are higher than the costs used for the scaling of the scoring function, then the agent's utility will always be negative.

Lemma 2. *If $c_t(\theta) > c_s(\theta)$ then $\bar{U}(\theta) < 0$ for any θ .*

Table 2. Comparison of Quadratic, Spherical and Logarithmic Scoring Rules

Scoring Rule:	Quadratic	Spherical	Logarithmic
θ^*	$\left(\frac{c_2}{c_1}\right)^2 \theta_0$	$\left(\frac{c_2}{c_1}\right)^{\frac{4}{3}} \theta_0$	$\left(\frac{c_2}{c_1}\right) \theta_0$
$\bar{P}(\theta_0)$	$c_2 \theta_0 (2 \frac{c_2}{c_1} - 1)$	$c_2 \theta_0 (4 \left(\frac{c_2}{c_1}\right)^{\frac{1}{3}} - 3)$	$c_2 \theta_0 (1 + \log\left(\frac{c_2}{c_1}\right))$

Note that costs are given by linear functions, $c(\theta) = c\theta$, and c_1 and c_2 are the lowest and second lowest costs.

Proof. Concavity of the expected score $\bar{S}(\theta)$ implies:

$$\bar{S}'(\theta_0)(\theta - \theta_0) \geq \bar{S}(\theta) - \bar{S}(\theta_0)$$

Similarly, convexity of the cost function $c_s(\theta)$ gives:

$$c_s'(\theta_0)(\theta - \theta_0) \leq c_s(\theta) - c_s(\theta_0).$$

By performing basic manipulations this results in:

$$\frac{c_s'(\theta_0)}{\bar{S}'(\theta_0)}(\bar{S}(\theta) - \bar{S}(\theta_0)) + c_s(\theta_0) - c_s(\theta) \leq 0$$

Furthermore, since $c_t(\theta) > c_s(\theta)$, the following holds, for any θ :

$$\bar{U}(\theta) = \frac{c_t'(\theta_0)}{\bar{S}'(\theta_0)}(\bar{S}(\theta) - \bar{S}(\theta_0)) + c_s(\theta_0) - c_t(\theta) < 0$$

□

Having presented these two key lemmas, we now proceed to prove the four economic properties of our mechanism.

Theorem 1. *Truthful revelation of agents' cost functions in the first stage of the mechanism is a weakly dominant strategy.*

Proof. We prove this by contradiction. Let $c_t(\theta)$ and $\hat{c}(\theta)$ denote an agents' true and reported cost functions respectively. Furthermore, let $c_s(\theta)$ denote the cost function used to scale the scoring function if the agent wins (i.e. if $\hat{c}(\theta_0) < c_s(\theta_0)$). Now, suppose that the agent misreports, but this does not affect whether the agent wins or not. If the agent loses then the payoff is always zero. If the agent wins the payoff is unaffected, since it is calculated from the second-lowest cost. Therefore, there is no incentive to misreport.

Suppose that the agent misreports, and now it does affect whether the agent wins or not. There are now two cases: (1) $c_t(\theta_0) > c_s(\theta_0)$ and $\hat{c}(\theta_0) < c_s(\theta_0)$ (the agent wins by misreporting but would have lost when truthful), and (2) $c_t(\theta_0) < c_s(\theta_0)$ and $\hat{c}(\theta_0) > c_s(\theta_0)$ (the agent loses by misreporting but would have won when truthful).

Case (1). Since the true cost $c_t(\theta_0) > c_s(\theta_0)$, it follows directly from lemma 2 that the expected utility $\bar{U}(\theta)$ 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 won by being truthful, but now receives a utility of zero. To show that this type of misreporting is suboptimal, we need to show that, when $c_t(\theta_0) < c_s(\theta_0)$, an agent benefits from being selected and generating the (optimal) estimate (i.e. $U(\theta^*) > 0$ when $c_t(\theta_0) < c_s(\theta_0)$). Now, since θ^* is optimal by definition, then $U(\theta^*) \geq U(\theta_0)$. From the expected utility in equation 6 we have, $\bar{U}(\theta_0) = c_s(\theta_0) - c_t(\theta_0) > 0$ when $c_t(\theta_0) < c_s(\theta_0)$, and hence $U(\theta^*) > 0$ at true costs reporting. □

Theorem 2. *The mechanism is incentive compatible w.r.t. the agent's reported measurement and precision in the second stage.*

Proof. The proof for this theorem follows directly from the definition of the strictly proper scoring rules (see section 3). □

Theorem 3. *The two-stage mechanism is individually rational.*

Proof. From theorem 1 we can assume that agents report their true cost functions in the first stage. Since agents who do not win in the first stage receive zero utility, we only need to consider the case of the selected agent with cost function $c_t(\theta) \leq c_s(\theta)$. From equation 6, it follows that $\bar{U}(\theta_0) = c_s(\theta_0) - c_t(\theta_0) \geq 0$. Lemma 1 shows that the agent may produce an estimate $\theta^* > \theta_0$. Since θ^* is optimal by definition, then $\bar{U}(\theta^*) \geq \bar{U}(\theta_0)$, and thus $\bar{U}(\theta^*) \geq 0$. □

Theorem 4. *For the agent selected in the first stage of the mechanism, it is optimal to produce an estimate with a precision equal or higher than the precision required by the centre, i.e., $\theta^* \geq \theta_0$.*

Proof. This proof follows directly from Lemma 1. In more detail, given that the agents reveal their true cost functions, we have $c_t(\theta) \leq c_s(\theta)$. Therefore, from lemma 1 it follows that $\theta^* \geq \theta_0$. □

Note that these proofs indicate that the two stages of the mechanism are inextricably linked and cannot be considered in isolation of one another. Indeed, apparently small changes to the second stage of the mechanism can destroy the incentive compatibility property of the first stage. For example, it is important to note that our mechanism is more precisely known as *interim* individually rational, since the utility is positive in expectation. In any specific instance, the payment could actually be negative if the prediction turns out to be far from the actual outcome. An alternative choice for the second stage of the mechanism would be to set β such that the payments are always positive, thus making the mechanism ex-post individually rational. However, this would then violate the incentive-compatibility property since the agents could then receive positive payoffs by misreporting their cost functions. Likewise, it might be tempting to imagine that the centre could use the revealed costs of the agents in order to request a lower precision, confident in the knowledge that the selected agent will actually produce an estimate of the required precision. However, by effectively using the lowest revealed cost within the payment rule in this way, the incentive-compatibility property of the mechanism would again be destroyed.

5 EMPIRICAL EVALUATION

Having proved the economic properties of the mechanism in the general case with any convex cost function, we now present empirical results for a specific scenario in which costs are linear functions, given by $c_i(\theta) = c_i\theta$, where the value of c_i is drawn from a uniform distribution $c_i \sim \mathcal{U}(1, 2)$ and $\theta_0 = 1$. Within this scenario our intention is to compare the performance of the three scoring presented earlier. To this end, for a range from 2 to 20 agents participating in the mechanism, we simulate the mechanism 10^6 times and, for each iteration, record the payment made to the agent who provided the estimate




Figure 1. The mean payment made by the centre.




Figure 2. The mean optimal precision of agents' estimates.

and the precision of this estimate. In figures 1 and 2 we present the means of these results (and note that the standard error in both means is much smaller than the symbol size).

Consider first figure 1 which shows the mean payment made by the centre. We note that, as expected, as the number of agents increases, the mean payment decreases toward the lower limit of the uniform distribution from which the costs were drawn. Furthermore, note that there is a fixed ordering over the entire range, with the payment resulting from the quadratic scoring rule being the highest, and that of the logarithmic scoring rule being the lowest.

In this figure, we also show the mean of the lowest and second lowest costs evaluated at the required precision θ_0 (denoted by $\bar{c}_1\theta_0$ and $\bar{c}_2\theta_0$ respectively). The first cost represents the minimum payment that could have been made if the costs of the agents were known to the centre. While, the second represents the payment that would have been made, had the agent produced an estimate of the required precision rather than its own optimal precision. The gap between $\bar{c}_1\theta_0$ and $\bar{c}_2\theta_0$ represents the 'information rent' that must be paid in the case that costs are unknown. The gap between $\bar{c}_2\theta_0$ and the mean payment of any particular scoring rule represents the loss that the centre has to cover due to the agent making a more precise estimate than required. The goal in selecting scoring rules is clearly to minimise this gap, and it can be seen that the logarithmic scoring rule is closest to achieving this goal. The reason for this can be seen in figure 2 where the precision of the estimates that were actually made are shown. Note that in this figure the logarithmic scoring rule is shown to induce agents to produce estimates closer to the required precision than both the spherical and the quadratic scoring rules. The same ordering as observed in these figures (when averaged over costs drawn from a uniform distribution) is also seen in analytical results for any

specific values for the lowest and second lowest costs (see table 2).

Based solely on these results, it can be considered that the logarithmic scoring rule presents the best choice for the centre in this case. However, it is important to note that the logarithmic scoring rule is unbounded. That is, in the event that the agent's estimate is far from the actual outcome, then a payment based on the logarithmic scoring rule will go to $-\infty$ since the agent's probability density function goes to 0 in this case (see row 1 of table 1). Thus, given this additional observation, it is clear that the spherical scoring rule represents a better choice since its payments are only slightly greater than that of the logarithmic, but it has finite bounds.

6 CONCLUSIONS

In this paper we introduced a novel two-stage mechanism based on strictly proper scoring rules that motivates selfish rational agents to make a costly probabilistic estimate or forecast of a specified precision and report it truthfully to a centre. We applied the mechanism in a setting in which the centre is faced with multiple agents but has no knowledge about their costs, and we proved that it was incentive compatible and individually rational. We also empirically evaluated our mechanism, and in comparing the quadratic, spherical and logarithmic scoring rules, showed that the logarithmic one minimises the centre's expected payment, but is unbounded. Thus, we proposed the use of the spherical rule as the best compromise between achieving minimal payments with finite bounds.

Our future work consists of two main tracks. First, we would like to explore the design of alternative strictly proper scoring rules, with the intention of minimising the loss that the centre has to cover, as a result of agents making an estimate of precision higher than the required one. In this respect the value of $\bar{c}_2\theta_0$, shown in figure 1, represents a bound on the ultimate performance of the mechanism. Second, we would like to extend our mechanism to the case where the centre procures estimates from more than one agent, and then fuses them together. When costs are convex, procuring several low precision estimates may be more cost effective than procuring a single high precision estimate. Indeed, Miller et al. have shown how scoring rules can be used to score one agent's estimate against another's, and thus in this case there is no need to wait until the actual event's outcome is revealed before making payments to agents [4]. However, in such a case, it is an open question as to whether it is possible to design a mechanism that incentivises multiple agents to truthfully reveal their costs and estimates.

ACKNOWLEDGEMENTS

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.

REFERENCES

- [1] A. D. Hendrickson and R. J. Buehler, 'Proper scores for probability forecasters', *The Annals of Mathematical Statistics*, **42**(6), 1916–1921, (1971).
- [2] R. Jurca and B. Faltings, 'Reputation-based service level agreements for web services', in *Proceedings of the International Conference on Service Oriented Computing (ICSOC)*, pp. 396–409, (2005).
- [3] J. E. Matheson and R. L. Winkler, 'Scoring rules for continuous probability distributions', *Management Science*, **22**(10), 1087–1096, (1976).
- [4] N. Miller, P. Resnick, and R. Zeckhauser, 'Eliciting honest feedback: The peer prediction method', *Management Science*, **51**(9), 1359–1373, (2005).
- [5] L. J. Savage, 'Elicitation of personal probabilities and expectations', *Journal of the American Statistical Association*, **66**(336), 783–801, (1977).
- [6] R. Selten, 'Axiomatic characterization of the quadratic scoring rule', *Experimental Economics*, **1**(1), 43–61, (1998).
- [7] A. Zohar and J. S. Rosenschein, 'Robust mechanisms for information elicitation', in *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 1202–1204, (2006).