

Eliciting Expert Advice in Service-Oriented Computing

Enrico H. Gerding¹, Kate Larson², and Nicholas R. Jennings¹

¹ University of Southampton, Southampton SO17 1BJ, United Kingdom
{eg, nrj}@ecs.soton.ac.uk

² University of Waterloo, Waterloo, ON N2L 3G1, Canada
klarson@cs.uwaterloo.ca

Abstract. We consider a setting where a consumer would like to elicit independent but costly reports from third-party experts about the reliability of a number of service providers. These reports can be of variable accuracy, but more accurate reports will be more costly to produce. The consumer can fuse reports from several experts to choose the provider with the highest probability of success. The goal in this paper is to find a mechanism which incentivises the experts to truthfully reveal the accuracy of the reports, and to induce the experts to invest costly resources in order to increase this accuracy. The challenge in doing so is that, while we can verify the success or failure of the selected service provider, we have no feedback about those service providers which were not selected. Moreover, we need to determine how to reward individual experts when the choice of service provider is based on a fused report from all experts. We explore a number of mechanisms to address this setting, including scoring rules, and indicate the problems in obtaining both truth telling and inducing the experts to produce accurate reports. We present a partial solution to this problem, and discuss remaining challenges.

1 Introduction

We consider the problem where a consumer requires a task to be completed, and can choose between one of several service providers to execute this task. We assume that these service providers are unreliable and the task may fail to be completed, but the consumer has no information about the reliability. Such a setting with execution uncertainty is becoming increasingly relevant, in particular within service-oriented computing where resources can be dynamically procured from third party suppliers, but also applies to more general settings where agents need to procure products and services of unknown quality. In this particular work we focus on obtaining trust and reputation information from expert agents who provide independent reports about the available service providers. To address this problem, we make use of techniques from the field of computational mechanism design to incentivise experts to produce accurate and truthful reports.

In more detail, we assume a setting with several expert agents, each of whom can obtain independent reports about the reliability of several service providers who are able to execute the relevant task. Based on these recommendations, the best service provider (in terms of their ability to execute the task) will then be selected. Crucially, however, the accuracy of the reports depends on the amount of effort (i.e. resources) that the

experts are willing to invest. For example, agencies that produce consumer advice often rely on reports from independent sources, where the accuracy of a report depends on the time and effort invested in testing and evaluation. Furthermore, some experts may have considerable experience with some providers, but not with others. Given this, we would like to design incentives to induce the agents to truthfully reveal their confidence in their reports, so that the consumer can combine them in such a way to enable an informed choice about which service provider to select. Furthermore, we would like to incentivise the experts to invest effort and in doing so increase the accuracy of their reports.

Our work is closely related to a number of papers that consider a mechanism design solution when faced with execution uncertainty, in particular [8, 2] consider the setting where service providers have a privately known probability of success, which needs to be elicited together with the cost of the service. In contrast, in our case the information about the probability of success is given by independent experts who do not benefit from a particular provider being selected. Indeed, we argue that having third party reviews, such as expert and user ratings, is very common in practice. Furthermore, we consider the setting where the reports may be costly, and where the experts not only reveal the probability of success, but also the number of ratings on which this report is based. As we argued above, this is especially important when fusing information from various sources into a single, combined recommendation, which forms the basis on which the service will be selected.

Information elicitation can also be achieved through the use of scoring rules. Originally introduced in statistical analysis, these are explicitly designed to evaluate probabilistic information, and we also study scoring rules as part of our mechanism in this paper. Typically, these rules produce payments which are conditioned on the outcome of an event, e.g. the success or failure of a task. Scoring rules have been used in related work to elicit information and induce effort, for example in [6, 9]. Our work differs from this and related approaches, however. First, in our case experts need to make observations about multiple service providers, but in the end only one of these providers is selected. This means that, if the scoring rule is conditioned on the outcome, only the information about the selected provider can be evaluated. The experts may therefore misreport their information about other providers. This can partially be addressed using a peer-prediction method, where a payment is not conditioned on the actual outcome, but on the reports submitted by other agents, but we show that this approach has other limitations in our setting.

Specifically, the contributions in this paper are as follows. First, we apply and compare a mechanism design approach as well as a scoring rule approach based on peer prediction. We show that neither of these approaches induce the agents to make accurate reports in our setting. To address this, we extend existing scoring rules so that agents are rewarded depending on the contribution that their reports have on the final outcome. However, we show that, under certain conditions, using this payment rule experts will have an incentive to misreport in order to try and manipulate the choice of the consumer. We argue that this effect can be mitigated by combining several approaches, but is not completely eliminated. Finally, we discuss some more challenges that remain to be addressed.

The remainder of the paper is structured as follows. We first present the problem and the desirable properties of the payment rules more formally in Section 2, and then apply the mechanism design approach to achieve truthful revelation in Section 3. We also show that the mechanism design approach does not induce effort and suggest instead the scoring rule approach in Section 4. In Section 5 we discuss some open problems, and conclude in Section 6.

2 The Problem Description

We assume that there exists a consumer C , with a single task, T , that it wants completed. The consumer derives a value V if the task is successful and zero otherwise. The consumer, however, is not able to complete the task itself, and instead must rely on a *service provider* to do the task for it. The consumer can choose from amongst a set of service providers, $SP = \{\emptyset, 1, 2, \dots, m\}$, each of which are capable of performing the task with the exception of \emptyset which represents the option of selecting no provider. If chosen, provider $k \in SP$ charges a fee ρ_k before attempting the task, and this fee is known to the consumer, C . While each service provider $k \in SP$ is capable of completing task T , not all are reliable and thus may fail before completing T . We let PoS_k denote the probability that service provider k will successfully complete the task, if asked to do so by C . While C would like to select the most reliable service provider possible, given the fee, PoS_k is private information, known only to the service provider. Finally, we assume that $PoS_\emptyset = 0, \rho_\emptyset = 0$. That is, C can always decide to not use a service provider (and thus get value 0 for not completing the task).

While C has no information about the reliability of the service providers, we assume that it can ask for information from a group of *experts*, $N = \{e_1, \dots, e_n\}$. Each agent $e_i \in N$ has some *experience* with each service provider. In particular, we assume that expert e_i has interacted with provider j in the past, and thus has observed $a_i(j)$ successes and $b_i(j)$ failures, where these successes and failures are drawn from the true underlying distribution of the service provider. Given these observations, an expert computes $n_i(j) = a_i(j) + b_i(j)$, the total number of observations made of service provider j , and $p_i(j) = a_i(j)/n_i(j)$, the observed probability of success of service provider j . We let $\omega_i = ((a_i(1), b_i(1)), \dots, (a_i(m), b_i(m)))$ denote all of e_i 's observations on all providers, and use the notation $\omega_i(j)$ to represent expert e_i 's observations of service provider j , $\omega = (\omega_1, \dots, \omega_n)$ and $\omega_{-i} = (\omega_1, \dots, \omega_{i-1}, \omega_{i+1}, \dots, \omega_n)$. Finally, we assume that it is costly for experts to gain expertise, and that each expert e_i incurs a cost $c_i \geq 0$ for each observation made, where c_i is public information.¹

Given the model just described, we want to design a *mechanism* so that the consumer can gather information from the individual experts, and, using the combined information, select the best service provider from amongst SP . That is, we would like the center to select:

$$k^* = \arg \max_{k \in SP} (V \cdot Pos_k(\hat{\omega}) - \rho_k) \quad (1)$$

¹ For simplicity, we assume that costs are linearly increasing with the number of observations and that costs are service-provider independent. Our results, however, apply to more general cost functions.

where $Pos_k(\hat{\omega})$ is the estimated probability that provider k will successfully complete the task, given the information provided to the consumer by the experts. If each e_i announces information $\hat{\omega}_i$ (where $\hat{\omega}_i$ may or may not be equal to ω_i , the true experience of e_i), $Pos_k(\hat{\omega})$ is computed as follows:

$$Pos_k(\hat{\omega}) = \frac{\sum_{i \in N} \hat{a}_i(k)}{\sum_{i \in N} (\hat{a}_i(k) + \hat{b}_i(k))}. \quad (2)$$

To make our model more concrete, consider the following example. Assume there are two service providers, A and B , and two experts, e_1 and e_2 . Expert e_1 has observed service provider A a total of twenty times, fifteen of which service provider A successfully completed the requested task. Expert e_1 only observed provider B 's attempt to execute a task once, on which the service provider failed. On the other hand, expert e_2 has a lot of experience with service provider B , having observed it succeed eighty times out of eighty-five attempts. However, e_2 has never interacted with service provider A and thus has no experience to report. Thus, $\omega_1 = ((15, 5), (0, 1))$ and $\omega_2 = ((0, 0), (80, 5))$. Assuming that both service providers have the same fee, ρ , and that the consumer knows ω_1 , and ω_2 , the consumer selects the service provider that maximizes $\max(V \cdot \frac{15}{20} - \rho, V \cdot \frac{80}{86} - \rho, 0)$. That is, the consumer would select provider 2.

The challenge faced by the consumer is two-fold. First, the consumer would like the reports made by the experts to be based on their true observations, i.e. that the mechanism is *truth-revealing*. In doing so, we consider mechanisms that incentivise truth-telling in *ex-post* Nash equilibrium.²

Definition 1 *A mechanism is truth-revealing (in ex-post Nash) if for all $e_i, e_j \in N, e_i \neq e_j$, when $\hat{\omega}_j = \omega_j$, reporting $\hat{\omega}_i = \omega_i$ maximises expert e_i 's expected utility.*

Second, the consumer would like to encourage the agents to gain as much experience as possible since having more observations of the service providers leads to a more accurate calculation of the true underlying Pos_k for each $k \in SP$.

Definition 2 *A mechanism is effort-inducing if experts receive a higher average payment by reporting a larger number of observations.*

In order to create the right incentives for the mechanism to be both truth-revealing and effort-inducing, the consumer, C , uses *transfers*. The transfer to expert $e_i \in N$, $\tau_i(\hat{\omega}|x) \in \mathbb{R}$, depends on the reported observations of *all* experts ($\hat{\omega}$), and on whether the service provider *chosen* by the consumer (i.e. k^*) is successful or not. That is $x \in \{\text{success, fail}\}$.

Assuming that $\tau_i(\hat{\omega}|x)$ is given, the utility for e_i when it observes ω_i , reports $\hat{\omega}_i$ and service provider k^* is selected is:

² We consider ex-post Nash instead of dominant strategies since experts have interdependent valuations, where the value of an agent depends on the chosen allocation (see Section 3 for more details), in which case generally no efficient, dominant strategy implementation is possible [3].

$$\begin{aligned}
U_i(\omega_i, \hat{\omega}_i) &= PoS_{k^*}(\omega_i) \tau_i(\hat{\omega}_i, \hat{\omega}_{-i} | success) \\
&\quad + (1 - PoS_{k^*}(\omega_i)) \tau_i(\hat{\omega}_i, \hat{\omega}_{-i} | fail) \\
&\quad - c_i \sum_{k \in SP} (a_i(k) + b_i(k)),
\end{aligned} \tag{3}$$

where

$$PoS_k(\omega_i) = \frac{a_i(k)}{a_i(k) + b_i(k)} \tag{4}$$

is the expert's private belief, based on its observations, of the true probability of success for service provider k . We note that the utility of e_i depends on several things, including the choice of service provider k^* , the reports of other experts since their reports influence the choice of k^* , and on the number of observations made.

Given transfer functions, we also derive the expected utility of the consumer:

$$\begin{aligned}
U_C(\hat{\omega}) &= -\rho_{k^*} + PoS_{k^*}(\hat{\omega}) \left(V - \sum_{i \in N} \tau_i(\hat{\omega} | success) \right) \\
&\quad - (1 - PoS_{k^*}(\hat{\omega})) \sum_{i \in N} \tau_i(\hat{\omega} | fail).
\end{aligned} \tag{5}$$

The rest of the paper is devoted to the study of different transfer functions, and the properties (i.e. truth-revealing and effort-inducing) they induce.

3 The Mechanism Design Approach

For many task-allocation problems, the Vickrey-Clarke-Groves (VCG) mechanism is used. In the VCG mechanism, agents are asked to reveal their private information (called their type), and given this information, the consumer (in our case) chooses an outcome which maximizes the social welfare, based on the information provided. In exchange, each agent is paid a transfer which is equal to their marginal contribution to the system. This mechanism is individually rational (i.e. no agents are harmed by participating) and truth-revealing.

Unfortunately, the VCG mechanism is not appropriate for our problem since it assumes that agents have *independent types*. Instead, our problem domain is one where agents have *interdependent types*, since the observations of a single agent only provides a partial model of the reliability of the service providers, and the actual service-provider model used by the consumer when making its choice is obtained by *fusing* the types of all agents.

While a series of papers have illustrated that it is impossible to design an incentive-compatible efficient (i.e. social-welfare maximizing) direct mechanism for settings where agents have interdependent types [1, 3], Mezzetti showed that it is possible to design incentive-compatible efficient mechanisms if the mechanism works over two stages [5]. First, the mechanism asks for agents' types and selects the outcome based on the reports. Then, after the outcome is realised, the agents report their actual payoffs. The transfers are computed using both the revealed types and the realised payoffs. This

mechanism is (weakly) truth-revealing in that agents are best-off revealing their types in the first stage, and are weakly best-off revealing their realised payoffs once the outcome has been selected.

We extend the framework described by Mezzetti to our setting.

Definition 3 (Two-Step Constant Mechanism) *Let N be the set of agents, and SP be the set of service providers. The two-step constant mechanism progresses as follows:*

1. Each $e_i \in N$ announces $\widehat{\omega}_i$.
2. The consumer selects service provider k^* where

$$k^* = \arg \max_{k \in SP} (V \cdot PoS_k(\widehat{\omega}) - \rho_k).$$

3. Service provider k^* attempts the task and either fails or succeeds.
4. Transfers to each $e_i \in N$ depend on k^* 's fee and on the success or failure of k^* ,

$$\tau_i(\widehat{\omega}|x) = \begin{cases} K - \rho_{k^*} & \text{if } x = \text{success} \\ -\rho_{k^*} & \text{if } x = \text{fail} \\ 0 & \text{if } k^* = \emptyset \end{cases}$$

where K is a constant chosen by the consumer so that $K - \rho_{k^*} > 0$.

Theorem 1 *The Two-Step Constant Mechanism is truth-revealing.*

Proof. Assume that all agents in N , but e_i , are reporting their observations truthfully. That is $\widehat{\omega}_{-i} = \omega_{-i}$. Since K is an arbitrarily defined constant, $K > 0$, if $k^* = \arg \max_{k \in SP} (V \cdot PoS_k(\widehat{\omega}) - \rho_k)$ then $k^* = \arg \max_{k \in SP} (K \cdot PoS_j(\widehat{\omega}) - \rho_k)$.

Assume that when reporting $\widehat{\omega}_i \neq \omega_i$, the consumer selects service provider $k' = \arg \max_{k \in SP} (V \cdot PoS_k((\widehat{\omega}_i, \omega_{-i})) - \rho_k)$ where $k' \neq k^*$, the service provider that would have been chosen if e_i revealed ω_i . Note also that by revealing $\widehat{\omega}_i$, agent e_i does not change the underlying probability that a particular service provider will fail, and that the best estimate that it will be successful is the fused estimate based on the actual observations of all agents, $PoS_k(\omega)$. Now,

$$\begin{aligned} U_i(\widehat{\omega}_i, \omega_{-i}) &= PoS_{k'}(\omega)(K - \rho_{k'}) + (1 - PoS_{k'}(\omega))(-\rho_{k'}) \\ &= PoS_{k'}(\omega)K - \rho_{k'} \\ &\leq PoS_{k^*}(\omega)K - \rho_{k^*} \\ &= U_i(\omega_i, \widehat{\omega}_{-i}). \end{aligned}$$

That is, if all other agents are reporting their true observations, then agent e_i is also best-off revealing its true observations. \square

While the mechanism just described is truth-revealing, it is not effort-inducing. All expert agents receive the same transfer, independent of the quality of the information they provide, even if the provided information was wildly inaccurate or uninformed. This is, in some sense, *unfair* and we would prefer to reward agents for the *quality* of their information. We believe that quality is directly related to *effort* in that the more effort an agent has made in collecting observations, the more reliable its information is. The Two-Step Constant Mechanism does not directly nor indirectly reward effort since it does not distinguish, in terms of transfers, between accurate and inaccurate agents.

4 The Scoring Rules Approach

As shown in the previous section, the mechanism design approach does not induce effort when the agents have costs for producing observations. In this section we consider alternative approaches where transfers are calculated using scoring rules, which are explicitly designed to elicit probabilistic information. We start by providing background on strictly proper scoring rules, and then propose two different transfer functions based on these rules. In our first approach, the *marginal-contribution scoring rule*, agents' transfers depend on the outcome (i.e. whether the chosen service provider successfully completed the task) as well as the impact that an agent's report had on the consumer's decision of service provider. In our second approach we consider a *peer-prediction* method where an agent's transfers depend only the reports of other agents, and not on the performance of the chosen service provider.

4.1 Background On Strictly Proper Scoring Rules

In this section we briefly provide background on scoring rules which were initially introduced in statistical analysis to evaluate expert probabilistic forecasts about some future event, e.g. a weather forecast, but are now increasingly being applied to agent-based applications [7, 9, 4]. In general, a scoring rule is a function which calculates a reward for an agent based on a distribution announced by the agent (it's prediction of some event) and the actual materialised event. A *strictly proper* scoring rule is a rule where an agent maximises its expected reward when the prediction it reveals is its actual belief about the event in question.

While there are many different strictly-proper scoring rules, in this paper we focus on the *logarithmic scoring rule*:

$$\begin{aligned} S(\hat{p}_i|success) &= \ln(\hat{p}_i), \\ S(\hat{p}_i|fail) &= \ln(1 - \hat{p}_i), \end{aligned} \tag{6}$$

where $\hat{p}_i \in [0, 1]$ is agent i 's reported belief that event $x = \text{success}$ will occur. Given this scoring rule, an agent's expected reward is $U_i(p, \hat{p}) = p \ln(\hat{p}) + (1 - p) \ln(1 - \hat{p})$, which is maximised when $p = \hat{p}$ (i.e. the agent is best-off revealing its true prediction of the event). We also note that if the function $S(p|x)$ is strictly proper, then so is the function $\alpha S(p|x) + \beta$, where $\alpha > 0$ and $\beta \in \mathbb{R}$ are scaling parameters.

Unfortunately, we are not able to directly use scoring rules, like the logarithmic scoring rules, directly as transfer functions for our service-provider problem. First, scoring rules, as just introduced, assume that probabilistic information for binary events (like the success or failure of a service provider) is captured by a single value, p , the probability of success. Thus, two agents who both reported $\hat{p} = \frac{1}{2}$ would be rewarded similarly, even though one agent may have only conducted two observations (and seen one success and one failure) while the other agent may have observed one hundred successes out of two hundred observations. Second, the standard scoring rules rely on observing the realisation of the event being predicted. In our setting, we want agents to report predictions on multiple service providers, only one of which is actually chosen and observed.

In order to deal with such situations, researchers have developed *peer-prediction methods* which rely on generating rewards based on how well the predictions of an agent agree with predictions of other agents [6]. We describe our proposed peer-prediction approach in Section 4.3.

4.2 Marginal Contribution Scoring Rule Based on Reality

In this section we introduce a scoring rule which rewards agents according to the *informativeness* of the information they provide, given the reports of other agents. Our *marginal-contribution* rule is given by:

$$\begin{aligned}\tau_i(\widehat{\omega}|success) &= \alpha_i \ln \left(\frac{PoS_{k^*}(\widehat{\omega})}{PoS_{k^*}(\widehat{\omega}_{-i})} \right), \\ \tau_i(\widehat{\omega}|fail) &= \alpha_i \ln \left(\frac{1 - PoS_{k^*}(\widehat{\omega})}{1 - PoS_{k^*}(\widehat{\omega}_{-i})} \right),\end{aligned}\tag{7}$$

if $k^* \neq \emptyset$, and $\tau_i = 0$ otherwise (if no provider is selected), where $\alpha_i > 0$ is a scaling parameter. The careful reader will note that this scoring rule is actually the relative entropy between the probability distribution determined by all agents reported observations, and the distribution when agent i 's observations are excluded. In particular, the marginal-contribution rule measures the inefficiency in using distribution $PoS_{k^*}(\widehat{\omega}_{-i})$ to make a prediction as to the success of service provider k^* compared to using distribution $PoS_{k^*}(\widehat{\omega})$. The marginal-contribution rule rewards experts whose reported observations led to an *increase* in the probability of success, if the task succeeds, and results in a negative reward if the task fails. The opposite holds for agents whose observation reports decreased the probability of success. The transfer, or reward, for an agent who provides no new information is zero.

In order to prove various properties that arise when using the marginal-contribution rule as a transfer function, we need to carefully define the utility of an agent. In particular, we need to specify an agent's beliefs about the observations made by other agents since the transfer that an agent i receives depends both on the service provider chosen and on the *total number of observations* made by other agents. If there are few observations, then the observations of agent i may be more informative and thus result in a higher transfer. However, if other agents have already reported many observations on a particular service provider, then additional observations may be less informative, resulting in lower rewards. In the following we will slightly abuse notation, and denote by $a_{-i}(k), b_{-i}(k), n_{-i}(k)$ the *aggregate* observations by all providers except i , and $\omega_{-i} = ((a_{-i}(1), b_{-i}(1)), \dots, (a_{-i}(m), b_{-i}(m)))$.

We initially assume that agent i knows $n_{-i}(k)$, the aggregate number of observations for each provider by all other agents. However, they do not know whether these observations represent successes or failures. Furthermore, we assume that the agent maintains a probability p , which is its belief that another agent observes a success. Typically, this belief will be based on an agent's own experience so far, in which case $p = p_i(k)$, but this assumption is not necessary for the proofs that follow. Given this, the probability of exactly $a_{-i}(k)$ successes and $b_{-i}(k) = (n_{-i}(k) - a_{-i}(k))$ failures is given by the binomial distribution:

$$\mathcal{B}(a_{-i}(k); n_{-i}(k), p) = \frac{n_{-i}(k)!}{a_{-i}(k)!b_{-i}(k)!} p^{a_{-i}(k)}(1-p)^{b_{-i}(k)}.$$

Ignoring agent costs, and provided that all agents $j \neq i$ are truthful, agent i 's expected utility when provider k is selected is then given by:

$$U_i(\omega_i, \hat{\omega}_i | k, n_{-i}) = \sum_{j=0}^{n_{-i}} \mathcal{B}(j; n_{-i}, p) U_i(\omega_i, \hat{\omega}_i | k, \omega_{-i} = (j+1, j-n_{-i}+1)), \quad (8)$$

where:

$$\begin{aligned} U_i(\omega_i, \hat{\omega}_i | k, \omega_{-i}) &= PoS_k(\omega_i, \omega_{-i}) \tau_i(\hat{\omega}_i, \omega_{-i} | success) \\ &\quad + (1 - PoS_k(\omega_i, \omega_{-i})) \tau_i(\hat{\omega}_i, \omega_{-i} | fail) \\ &= PoS_k(\omega_i, \omega_{-i}) \alpha_i \ln \left(\frac{PoS_k(\hat{\omega}_i, \omega_{-i})}{PoS_k(\omega_{-i})} \right) \\ &\quad + (1 - PoS_k(\omega_i, \omega_{-i})) \alpha_i \ln \left(\frac{1 - PoS_k(\hat{\omega}_i, \omega_{-i})}{1 - PoS_k(\omega_{-i})} \right), \quad (9) \end{aligned}$$

where $PoS_k(\omega_i, \omega_{-i}) = (a_i + a_{-i}) / (n_i + n_{-i})$ is agent i 's beliefs about the probability of success when observing a_i out of n_i successes, conditional on other agents observing a_{-i} out of n_{-i} successes. Furthermore, $PoS_k(\omega_{-i}) = a_{-i} / n_{-i}$.³

Theorem 2 For chosen service provider, k , Equation 7 is a strictly proper scoring rule. That is, for $\omega \neq \hat{\omega}$:

$$U_i(\omega_i, \omega_i | k, n_{-i}) > U_i(\omega_i, \hat{\omega}_i | k, n_{-i})$$

Proof. Since $PoS_k(\omega_{-i})$ is a constant, we can ignore this part of the scoring rule. By doing so, note that Equation 9 can be written as:

$$U(p, q) = p \ln(q) + (1-p) \ln(1-q)$$

Now, from Gibbs' inequality it follows that $U(p, p) \geq U(p, q)$. Furthermore, $U(p, p) = U(p, q)$ if and only if $p = q$. This means that Equation 9 is maximised if and only if $PoS_k(\omega_i, \omega_{-i}) = PoS_k(\hat{\omega}_i, \omega_{-i})$, and therefore $\hat{\omega}_i = \omega_i$ is optimal. \square

We next show that the marginal-contribution scoring rule induces *effort*. In particular, we show that when $c_i = 0$ and agents are truthful, then the more observations an agent makes, the higher its expected utility.

Theorem 3 The scoring rule given by Equation 7 induces effort, that is, for any scalar $x > 1$, where $x \in \mathbb{N}^+$:

$$U_i(x\omega_i, x\omega_i | k, n_{-i}) > U_i(\omega_i, \omega_i | k, n_{-i})$$

³ While Equation 9 is not well defined for $a_{-i} = 0$ and $a_{-i} = n_{-i}$, we can assume that all agents use Laplace smoothing, thus avoiding these cases.

Proof. We prove this by showing that $dU_i(x\omega_i, x\omega_i|k, n_{-i})/dx > 0$ for any $x > 0, n_i > 0, a_{-i} \geq 1, b_{-i} \geq 1$. From Equation 8 we have:

$$\frac{dU_i(x\omega_i, x\omega_i|k, n_{-i})}{dx} = \sum_{j=0}^{n_{-i}} \mathcal{B}(a_{-i}; n_{-i}, p_i) \frac{dU_i(x\omega_i, x\omega_i|k, \omega_{-i})}{dx}$$

It is therefore sufficient to show that $dU_i(x\omega_i, x\omega_i|k, \omega_{-i})/dx \geq 0$ for all $a_{-i} \leq n_{-i}$, and that this inequality is *strict* for at least one $a_{-i} \leq n_{-i}$. To show that this holds, note that we can write the first derivative as follows:

$$\frac{dU_i(x\omega_i, x\omega_i|k, \omega_{-i})}{dx} = -\alpha \frac{b_i a_{-i} - b_{-i} a_i}{(x n_i + n_{-i})^2} \cdot \ln \left(\frac{b_{-i}(x a_i + a_{-i})}{a_{-i}(x b_i + b_{-i})} \right) \quad (10)$$

Since $\ln(x) \leq 1 - x$ and $\alpha > 0$, it is sufficient to show that:

$$-\frac{b_i a_{-i} - b_{-i} a_i}{(x n_i + n_{-i})^2} \cdot \left(\frac{b_{-i}(x a_i + a_{-i})}{a_{-i}(x b_i + b_{-i})} - 1 \right) \geq 0,$$

which simplifies to:

$$\frac{x(b_i a_{-i} - b_{-i} a_i)^2}{a_{-i}(x b_i + b_{-i})(x n_i + n_{-i})^2} \geq 0$$

Now, clearly the above inequality is equal to zero when $b_i = b_{-i}, a_i = b_{-i}$ (given $x > 0, n_i > 0, a_{-i} \geq 1, b_{-i} \geq 1$), but is strictly positive in all other cases. Hence it follows that $dU_i(x\omega_i, x\omega_i|k, n_{-i})/dx > 0$. \square

Clearly, if the utility increases as a function of x , then depending on an agent's cost c_i for producing an observation, the agent will be incentivised to produce a greater number of observations. We can set the desired level of effort by multiplying the transfer by an appropriately chosen parameter α . Since affine transformations of strictly proper scoring rules are also strictly proper, this does not change the properties of the transfers.

We now show that, despite Equation 8 being strictly proper, when $k = k^*$ is given by Equation 1, the mechanism as a whole is no longer truth-telling. In particular, the strictly-properness of the transfer function only holds for the observations made on the chosen service-provider, k^* . Agents may have incentive to misreport their information in order to affect the choice of service provider.

Theorem 4 *If k^* is given by Equation 1, and transfers τ_i for $i \in N$ are calculated according to Equation 7, then there exist settings in which an agent's expected utility is maximised for $\hat{\omega}_i \neq \omega_i$.*

Proof. The proof is by example. Suppose that there are two service providers, $k = 1, k = 2$ and $\rho_1 = \rho_2 = 0$ (fees are zero). Furthermore, suppose that for some agent $i \in N$, the following holds:

$$U_i(\omega, \hat{\omega}_i = \omega_i|k = 1) > U_i(\omega, \hat{\omega}_i = \omega_i|k = 2)$$

At the same time:

$$PoS_1(\omega) < PoS_2(\omega)$$

That is, the consumer prefers service provider 2, but agent i is better off if provider 1 is selected, e.g. because it has relatively many observations for that particular service provider. In this case, the agent could try and manipulate the outcome of k^* by mis-reporting. This can be done by either increasing the probability of success of provider 1, $\hat{p}_i(1)$ (by increasing $\hat{a}_i(1)$ or decreasing $\hat{b}_i(1)$), or decreasing $\hat{p}_i(2)$ (by decreasing $\hat{a}_i(2)$ or increasing $\hat{b}_i(2)$). Note that the first type of manipulation negatively affects $U_i(\cdot|k=1)$ and thus the utility of agent i if its favourite outcome is selected, but the second option does not. Therefore, an agent can always report $\hat{a}_i(2) = 0, \hat{b}_i(2) = \infty$ to ensure that option 1 is selected by the consumer. \square

4.3 Scoring Rule Based on Peer Prediction

In the previous section we introduced the marginal contribution mechanism and proved that it is effort-inducing, but not necessarily truth-telling since agents may try to manipulate their reports in order to influence the choice of the consumer. Since the transfers of the agents were based only on the outcome of the chosen service provider, lying about observations of non-selected service providers was not detected nor punished by the mechanism.

In this section we propose a *peer-prediction* based method, where agents' transfers are determined by comparing their announced observations with the reported observations of other agents. Since agents report observations made on the same set of service providers, their reports should be correlated, and peer-prediction methods try to detect unexpected differences in reported observations in order to detect mis-reporting agents.

Existing approaches for peer prediction calculate the score of an agent based on the report of a single other agent, called the reference agent [6]. Although this provides the necessary properties, in practice the reference agent may have few observations and the resulting score will be highly irregular. To address this, we use a virtual agent which contains the fused observations of all other agents. The transfers are calculated separately for reports about each provider k , and are given by:

$$\tau_i^k(\hat{\omega}_i|\hat{\omega}_{-i}) = \alpha \text{PoS}_k(\hat{\omega}_{-i}) \ln(\text{PoS}_k(\hat{\omega}_i)) + \alpha(1 - \text{PoS}_k(\hat{\omega}_{-i})) \ln(1 - \text{PoS}_k(\hat{\omega}_i)) + \beta \quad (11)$$

We now formulate an expert's expected utility. As before, we assume that the agents only know the total number of observations made by other agents with service provider k , $n_{-i}(k)$, and p , the beliefs about other agents observing a success. As opposed to the marginal contribution approach, now the agents receive a separate transfer for the observations of each service provider $k \in SP$. That is:

$$U_i(\omega_i, \hat{\omega}_i | n_{-i}) = \sum_{k \in SP} U_i^k(\omega_i, \hat{\omega}_i | n_{-i}), \quad (12)$$

where:

$$U_i^k(\omega_i, \hat{\omega}_i | n_{-i}) = \sum_{j=0}^{n_{-i}} \mathcal{B}(j; n_{-i}, p) \tau_i^k(\hat{\omega}_i | \omega_i(k) = (j, n_{-i} - j)). \quad (13)$$

Theorem 5 *The transfers defined in Equation 11 result in agents truthfully revealing their estimated probability of success for each service provider.*

Proof. We now show that truth telling maximises expected utility, and that this is strict in terms of reporting the probability of success (as opposed to number of successes and failures). Let $U_i^k(p_i, \hat{p}_i | n_{-i})$ denote the expected utility in terms of the probability of success, where $\hat{p}_i = \hat{a}_i / \hat{n}_i$. Clearly,

$$U_i^k(p_i, \hat{p}_i | n_{-i}) = U_i^k(\omega_i, \hat{\omega}_i | n_{-i})$$

Taking the first derivative results in:

$$\frac{dU_i^k(p_i, \hat{p}_i | n_{-i})}{d\hat{p}_i} = \frac{\alpha}{\hat{p}_i(1 - \hat{p}_i)} \sum_{j=0}^{n_{-i}} B(j; n_{-i}, p_i) \left(\frac{j}{n_{-i}} - \hat{p}_i \right) = \frac{\alpha(p_i - \hat{p}_i)}{\hat{p}_i(1 - \hat{p}_i)}$$

Clearly the expected utility is maximised if and only if $p_i = \hat{p}_i$. \square

While our peer-prediction method encourages all agents to truthfully reveal their probabilistic estimates about all service providers, it is not able to induce effort. In particular, *any* combination of observations resulting in the same probabilistic estimate results in the same utility for the agent. For example, an agent observing one success and one failure for a particular service provider is given the same reward as an agent who observes fifty successes and fifty failures.

5 Challenges

In the previous sections we described three different transfer functions that could be used by a mechanism designer in order to elicit experience information from experts. While each transfer proposal illustrated some desirable properties, we argue that none balanced the right combination of eliciting both truthful reports from the experts concerning their experience with different service providers as well as encouraging the experts to conduct many samples/experiments. In this section we discuss some of the challenges which we still face.

5.1 Balancing Truth-Telling and Effort

As we have seen, of the three rules described in this paper, only the marginal contribution scoring rule induces effort in our setting. However, using this scoring rule, an expert achieves a higher expected utility if an outcome is selected for which the expert has more observations (or the other agents have fewer observations) since this will, on average, increase the information gain and thus the score. Furthermore, an expert is scored based only on the result of the selected outcome, and its remaining observations are not rewarded. As a result, an expert may have an incentive to manipulate the outcome so that a suboptimal service provider k' is chosen, but for which he holds more observations, e.g. by making other outcomes less attractive. We note, however,

that this is only a problem when there is a relatively large imbalance in the number of observations from each provider.

A possible solution to the above problem is to produce a transfer which combines the marginal contribution scoring rule with one of the other two approaches. We discuss each of the two combinations in more detail, starting with the peer prediction method. The peer prediction approach provides a score for probabilistic estimates from each of the providers, independent of the selected provider and of the outcome. However, in our setting with discrete outcomes, this approach does not reward effort. That is, a single observation provides the same expected reward as having more than one. By having a linear combination of the two scoring rules, however, and by scaling them appropriately, the incentive to misreport the probability of success can be eliminated. Nevertheless, this approach is still not foolproof, as an expert can sometimes still profit by manipulating the outcome by simply reducing the reported variance (by proportionally increasing the number of reported successes and failures) whilst keeping the probability of success the same. As we have seen, this will not affect the score obtained through peer prediction. Although not perfect, however, the combination does reduce the number of cases in which a deviation from truth telling is profitable.

Alternatively, we can combine the marginal contribution scoring rule with the transfers from the mechanism design approach, again using a linear combination. The latter aligns the payoffs with the consumer's utility and therefore it is in the interest of the experts that the optimal provider k^* is selected. However, in the case that there is very little difference between the expected probability of success of different providers, it may still be profitable for an agent to try and manipulate the outcome in favour of a provider k' , if the agent has relatively more information about k' . In such a case, however, since the providers are very similar, if an agent misreports, this will have very little effect on the consumer's utility (since otherwise this would also negatively affect the agent because a significant part of the utility has been aligned with the consumer's utility).

5.2 The Expert's Decision Problem: What and How Many Observations to Make

So far we have mainly focused on the mechanism design problem, and we have ignored an important part of the expert's decision problem: Given that an effort-inducing mechanism with scaling parameter α , and given the costs, how many observations should an expert make to maximise its expected utility, if any? Moreover, which of the providers should the expert sample from? This is a complex problem which we will not address in this paper. Rather, we will briefly address some of the issues.

First, the expert's expected utility needs to be adequately modeled. As we already mentioned in Section 4.2, the expert needs to reason about the observations made by other experts since this will affect the information gain and therefore the reward obtained. In order to prove the various properties, in Section 4.3 we assumed that the agents knew *the number* of observations of other agents (but not whether these were successes or failures). In practice, however, this information may not be available since each agent is waiting to see what other agents will do before being able to make a decision. Essentially, since the optimal decision depends on the actions of other agents,

this is a strategic problem and can be addressed game-theoretically using the notion of Nash equilibrium.

Furthermore, even if an agent is able to determine the number of observations of others, in the case of the marginal contribution scoring rule, only those observations from the selected provider will be rewarded. An agent must therefore also reason about which provider is most likely to be selected by the consumer. However, it needs to sample all providers in order to reason about which provider is most likely to be the most successful on average. This problem is well studied in the literature and is referred as the multi-armed bandit problem. Although solutions exist, these will need to be adapted to this specific setting.

6 Conclusion

To conclude, we considered a setting where a consumer needs to choose a service provider to complete a task, and approaches several independent experts for their recommendations. We presented three mechanisms for eliciting this information from the experts: (1) a two-stage mechanism design approach based on the work by Mezzetti [5], (2) a scoring rule based on peer prediction, and (3) a scoring rule based on KL divergence. We showed that the first two approaches induce truth telling but do not induce effort if the information is costly to produce. The third approach does induce effort, but in some settings the experts gain by misreporting about the reliability of the service provider for which they hold no or little information. Although not foolproof, we argue that a combination of approaches mitigates most of the concerns and works in most cases. At the same time, there are a number of open problems which remain to be studied. In particular, experts need to decide which producers to evaluate, and how accurate the reports should be, given the costs and the mechanism. How to do this optimally is a challenging problem and remains for future work.

Acknowledgement

We would like to thank the reviewers for their useful comments. This research was undertaken as part of the EPSRC (Engineering and Physical Research Council) funded project on Market-Based Control (GR/T10664/01).

References

1. P. Dasgupta and E. Maskin. Efficient auctions. *Quarterly Journal of Economics*, 115:341–389, 2000.
2. E. H. Gerding, A. Rogers, K. Larson, and N. R. Jennings. Mechanism design for task procurement with flexible quality of service. In *Proc. of 8th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2009)*, 2009.
3. P. Jehiel and B. Moldovanu. Efficient design with interdependent valuations. *Econometrica*, 69:1237–1259, 2001.

4. R. Jurca and B. Faltings. Incentives for expressing opinions in online polls. In *Proceedings of the 9th ACM Conference on Electronic Commerce*, pages 119–128, 2008.
5. C. Mezzetti. Mechanism design with interdependent valuations: Efficiency. *Econometrica*, 72(5):1617–1626, 2004.
6. N. Miller, P. Resnick, and R. Zeckhauser. Eliciting informative feedback: The peer-prediction method. *Management Science*, 51(9):1359, 2005.
7. A. Papakonstantinou, A. Rogers, E. H. Gerding, and N. R. Jennings. A truthful two-stage mechanism for eliciting probabilistic estimates with unknown costs. In *Proc. of the 18th European Conference on Artificial Intelligence, Patras, Greece*, pages 448–452, 2008.
8. R. Porter, A. Ronen, Y. Shoham, and M. Tennenholtz. Fault tolerant mechanism design. *Artificial Intelligence*, 172(15):1783–1799, 2008.
9. A. Zohar and J. S. Rosenschein. Mechanisms for information elicitation. *Artificial Intelligence*, 172(16–17):1917–1939, 2008.