



Multi-dimensional procurement auction under uncertain and asymmetric information

Papakonstantinou, Athanasios; Bogetoft, Peter

Published in:
European Journal of Operational Research

Link to article, DOI:
[10.1016/j.ejor.2016.09.060](https://doi.org/10.1016/j.ejor.2016.09.060)

Publication date:
2017

Document Version
Peer reviewed version

[Link back to DTU Orbit](#)

Citation (APA):
Papakonstantinou, A., & Bogetoft, P. (2017). Multi-dimensional procurement auction under uncertain and asymmetric information. *European Journal of Operational Research*, 258(3), 1171–1180.
<https://doi.org/10.1016/j.ejor.2016.09.060>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Multi-dimensional procurement auction under uncertain and asymmetric information

Athanasios Papakonstantinou*

*Technical University of Denmark, Department of Electrical Engineering, Elektrovej 325,
DK-2800 Kogens Lyngby, Denmark*

Peter Bogetoft

*Copenhagen Business School, Department of Economics, Porcelaenshaven 16 A, DK-2000
Frederiksberg, Denmark*

Abstract

This article addresses two important issues in public procurement: ex ante uncertainty about the participating agents' qualities and costs and their strategic behaviour. We present a novel multi-dimensional auction that incentivises agents to make a partial inquiry into the procured task and to honestly report quality-cost probabilistic estimates based on which the principal can choose the agent that offers the best value for money. The mechanism extends second score auction design to settings where the quality is uncertain and it provides incentives to both collect information and deliver desired qualities.

Keywords: Auctions/bidding; Information asymmetry; Uncertainty modelling; Quality and cost uncertainty; Strictly proper scoring rules;

1. Introduction

Auctions are as popular as ever. Nowadays, auctions impact our everyday lives in direct, as well as in less obvious ways: from electronic trade, on-line marketplaces and Internet based business services, to on-line advertising and

*This research was *mainly* funded by the Center for research in the Foundations of Electronic Markets (CFEM) supported by the Danish Council on Strategic Research.

Corresponding author. Email address athpapa@elektro.dtu.dk (A. Papakonstantinou)

web-search engines. Such new services require new developments in auction theory.

Traditional auction theory has been focusing on single-dimensional price auctions that allocate an item to the highest bidder, or procure a service from the cheapest supplier. Although such auctions have been widely used to procure pre-specified services from suppliers at the lowest cost (Vickrey, 1961; Groves, 1973; Clarke, 1971), they may not be adequate to capture the complexities of the information based digital economy where the traded good is often not a well defined item. For example, in the purchase of a web-based service (e.g on-line storage, web hosting) parameters such as the available bandwidth, the robustness of the service, and the responsiveness may be just as important as the price of the service. Similarly, in government procurement, projects such as the construction of public infrastructure cannot be allocated solely to the cheapest contractor. Parameters such as the quality of the materials, the design and impact of the project on local communities and the environment often differ from one proposal to another and should be taken into consideration. Examples of such practices can be found on both sides of the Atlantic: the European Union implements a procurement directive calling for both lowest cost and best economic value (Asker & Cantillon, 2008), while the US Department of Defence requires that procurement competitions consider also performance and quality dimensions (Che, 1993).

In his seminal paper Che (1993), Che designed a series of multi-dimensional auctions (first score, second score and second preferred score) to address cases where not only the cost, but also the the quality of a product is important. In these auctions, suppliers report their production quality and the associated costs, and the mechanism maps the multi-dimensional bid into a single-dimensional evaluation referred as a ‘score’. All three auctions are incentive compatible, and based on the assumption that costs are independently distributed the first and second score auctions implement the socially optimal (allocatively efficient) outcome. Branco Branco (2007) relaxed the assumption of independent costs by introducing a two-stage multi-dimensional auction for

settings with correlation among suppliers' costs. A mechanism proposed by Bogetoft & Nielsen (2008) further exploited the correlations among the costs of different agents via Data Envelopment Analysis (Charnes et al., 1979, 1987) based competition.

More applications of score auctions can be found in Computer Science literature and in particular in multi-agent systems and e-commerce (He et al., 2003). Bichler (2000) paves the way for possible e-commerce applications of multi-dimensional auctions by showing that they result in a significantly higher utility when compared to single-dimensional auctions in a web-based experimental setting. Furthermore, Beil and Wein in Beil & Wein (2003) propose an iterative mechanism in which the principal sequentially estimates each bidder's cost function through a series of score auctions. Parkes and Kalagnanam (2005) also propose an iterative multi-attribute price-based procurement auction in which agents in each round submit their bids and a winner maximising the principal's preference is selected. They show that their mechanism terminates with a modified Vickrey-Clarke-Groves allocation. Multi-dimensional auctions can also be applied in settings where multiple suppliers are necessary to satisfy the principal's demand (Bichler & Kalagnanam, 2003).

Now, although these approaches address important issues in procurement, they rely on two significant assumptions. First, they assume that the principal will enforce truthful reporting of agents' qualities through external means. They do not give agents control of their reported (and consequently contracted) qualities, or they assume that in case the observed quality deviates from the reported one, the auction may be cancelled or heavy fines may be issued to the winner of the auction (Che, 1993). Second, existing literature assumes stochasticity solely for the production costs. In the very few cases where misreporting quality is considered (Papakonstantinou & Bogetoft, 2013), qualities are deterministic. We find this assumption particularly un-realistic, since often the agents qualities may depend on stochastic factors (i.e. availability of certain materials, weather conditions, faults). In such cases, the principal faces the significant challenge of

eliciting accurate probabilistic estimates of the agents’ qualities and based on them making a decision.

To address situations where it is costly to generate the aforementioned quality estimates, we introduce *strictly proper scoring rule* payment in a multi-dimensional auction. Strictly proper scoring rules are designed to elicit accurate predictions by rewarding forecasters based on how close the actual outcome is to their prediction (Savage, 1977; Hendrickson & Buehler, 1971; Friedman, 1983). Strictly proper scoring rules have been widely used in mechanism design to elicit accurate information and in particular to design reputation systems that promote truthful reporting on a service experience (Jurca & Faltings, 2005, 2006, 2007). Furthermore, it has been shown how an appropriately scaled strictly proper scoring rule can be used to incentivise agents to invest costly resources when generating their forecasts Miller et al. (2005, 2007). Extensions are given in Papakonstantinou et al. (2011) and Witkowski & Parkes (2012), with a brief summary of the main insights provided in Section 3.

In this paper we combine elements from multi-dimensional auctions and information elicitation mechanisms. Specifically, we consider a procurement setting where agents commit resources to estimate their production qualities, which they, in turn, report to the principal. These quality estimates are costly for all agents irrespective of the auction’s final outcome. This relates our work to all-pay auctions, where agents’ bids model the effort exerted in acquiring a ‘prize’. For example, Anderson et al. (1998) introduces an all-pay auction whereby the contract is allocated to the agent with the highest bid and the agents face bidding costs directly proportional to the actual bid. McAdams in McAdams (2006) considers a more general multi-unit multi-dimensional all-pay auction where agents’ costs are quadratic functions of their bids. The fact that in our proposed mechanism the principal allocates the task based solely on the agents’ estimates of their quality, which are not drawn from the same underlying distribution, differentiates our work from both aforementioned cases.

We propose a procurement auction where instead of having to complete the allocated task *ex ante* in order to participate, heterogeneous agents provide

estimates of the quality they intend to produce. The selected agent is allocated the project and once it is completed the principal observes the outcome and issues a payment that compensates for both data collection and production costs. As opposed to existing implementations of all-pay auctions where agents competing for a contract risk not getting paid even if they complete the task (cf. Liu et al. (2014)), in the proposed mechanism, the principal can penalise the selected agent depending on potential deviations between the agent’s reported estimated quality and the actual produced one.

We provide a solid theoretical foundation for our procedure. We show incentive compatibility and individual rationality. We also show that in expectation and under the use of a consistent set of Bayesian estimators, our mechanism achieves the outcome of the second score auction in which agents are able to directly report their actual quality outcomes. Finally, we evaluate our mechanism numerically in a setting where agents’ investigations are represented by samples of Gaussian distributions. The mechanism’s computational aspects are discussed and its convergence to the outcome of the second score auction is demonstrated through numerical simulations. The simulations complement our theoretical findings.

The rest of the paper is organised as follows: In Section 2 we describe the setting in more details, and in Section 3 we provide the relevant background in strictly proper scoring rules. In Section 4 we define the mechanism, while in Section 5 we outline and prove the economic properties. In Section 6 we evaluate the mechanism through numerical simulations and in Section 7 we conclude and provide some insights on future work.

2. The context

We consider a principal interested in procuring a task or a service from one of N rational and risk neutral agents. The provided task or service may be an independent task or part of a more complex one, without this affecting our analysis. The task can be characterised by multiple parameters in an s -

dimensional vector of qualities $y_0^i \in \mathbb{R}^s$ with $y_0^i > 0$ and $i \in I = \{1, \dots, N\}$. To simplify the analysis, however, we assume that for each agent the parameters of its service can be aggregated in one variable, $y_0^i \in \mathbb{R}$, with each agent having a single quality profile.

We depart from the existing literature by introducing uncertainty regarding the agents' qualities. We model uncertainty by assuming that each agent i commits variable degree of resources when estimating his individual quality y_0^i , which denotes the ground truth estimated by agent i but observed by the principal after the task is complete and if agent i was actually the selected provider. Let an agent's estimate be a predictive distribution denoted by a parametric distribution with mean y^i and precision θ^i i.e., $\mathcal{F}(y^i, \theta^i)$. The estimates are assumed to be unbiased such that y^i is independently drawn from $\mathcal{F}(y_0^i, \theta^i)$ i.e., $y^i \sim \mathcal{F}(y_0^i, \theta^i)$. We do not restrict our analysis to a specific distribution, although for the numerical evaluation of the mechanism we use Gaussian distributions following the data fusion literature Gregory (2005); DeGroot & Schervish (2002).

Under this information uncertainty framework, the estimate's precision represents the resources invested in generating that estimate. Hence, the cost associated with the collection of the information about the quality, increases as precision θ^i increases. We will therefore model data collection cost c^i as a convex, increasing and double differentiable function such as $c^i(\theta^i) = C^i \theta^{i^2}$, where $C^i > 0$ is a parameter which represents different base costs for each agent. We assume that the cost parameter is common knowledge, while the agents' precisions are private to each agent.

Now, regarding the production costs, we follow the existing literature (Che, 1993) by assuming that agents are capable of producing different levels of outputs, and that in order to produce the quality y_0^i agent i needs inputs which depend on each agent's efficiencies. These inputs are the costs involved in production and should not be confused with the costs involved in the estimation of the quality. Here, costs are private information to each agent and cannot be verified by any third party. The cost x^i an agent faces in the production of his

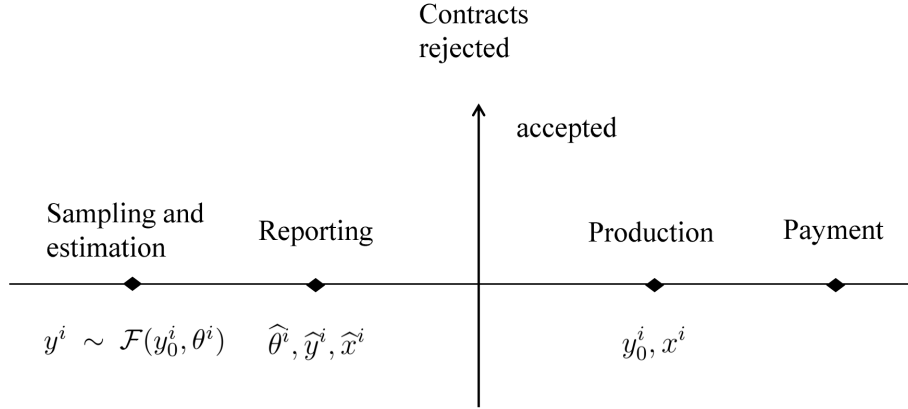


Figure 1: The time-line of the game.

quality is denoted as $x^i(y_0^i, l^i)$, where l^i represents the agent's private information about his production cost (in)efficiency. While agents are aware of their cost parameters, the principal has only access to their distribution. We assume that l^i is independently and identically distributed over $[\underline{l}, \bar{l}]$ with $0 < \underline{l} < \bar{l} < +\infty$ according to a distribution with positive and continuously differentiable density function. Finally, the cost function is increasing in both quality and the cost (in)efficiency parameter and it is convex in the quality.

Based on the above, the time-line (Figure 1) of the game is as follows. Initially, each agent collects information about his likely production quality and calculates the corresponding production costs. By sampling with precision θ^i , and spending information collection costs $c^i(\theta^i)$ he is able to predict his quality y^i and based on that predict the cost of the production as well. We assume that the agent can send possibly manipulated signals about his production quality level, his production costs, and the precision of his prediction to the principal before the principal decides on the provider. Let the signalled production be \hat{y}^i , the signalled data collection effort be $\hat{\theta}^i$, and the signalled cost be \hat{x}^i . The principal can use these signals to choose the provider and based on the realised quality y_0^i determine a reimbursement. If the principal picks agent i as the

provider, his value of the realised quality y_0^i will be given by $V(y_0^i)$ where $V(\cdot)$ is an increasing, concave and twice differentiable function of the quality.

To sum up, in this setting the principal has to deal with poor quality and cost estimates generated by agents not committing significant resources to the pre-bidding information collection, with misreporting of the estimates and with incentivising the winning agent to actually produce the final outputs.

3. Strictly proper scoring rules

Before turning to the details of the mechanism, it is convenient to discuss the simpler problem of inducing agents to collect information about their qualities and to reveal their findings.

So-called strictly proper scoring rules are used as a tool for eliciting forecasters' beliefs of future events in various domains ranging from meteorology and weather forecasting to computer science and on-line trust and reputation systems. Such scoring rules incentivise a risk neutral forecaster to truthfully report his forecast by maximising his expected reward. Imagine that a forecaster's prediction y of an event y_0 follows the distribution \mathcal{F} with the report denoted as $\hat{\mathcal{F}}$. The forecaster's corresponding score is then $S(y_0|\hat{\mathcal{F}})$ and the expected score is

$$\bar{S}(\hat{\mathcal{F}}, \mathcal{F}) = \int_{-\infty}^{\infty} \mathcal{F} S(y_0|\hat{\mathcal{F}}) dy_0 \quad (1)$$

The scoring rule S defined as *strictly proper* if its expected value, $\bar{S}(\hat{\mathcal{F}}, \mathcal{F})$, is maximised by truthful reporting i.e. $\bar{S}(\mathcal{F}, \mathcal{F}) \geq \bar{S}(\hat{\mathcal{F}}, \mathcal{F})$. Due to this property, a payment based on such a scoring rule can create incentives for truthful reporting.

Furthermore, strictly proper scoring rules can guarantee sufficient data collection effort on behalf of the agents. This process is described by Miller et al. (2005) who note that by using an affine function $\alpha + \beta S$ of a strictly proper scoring rule S as a payment, it is possible to induce an agent to make and truthfully report an estimate at a specific precision θ_0 .

In this case, the agent's expected payment, $\bar{P}(\theta)$, is

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

where α and β are the scaling parameters in the affine transformation, θ is the agent's true precision, and $\bar{S}(\theta)$ is the twice differentiable expected score function. Parameter α guarantees that the estimate will be generated at the appropriate precision, while β compensates the agent for the cost of his estimate. Consequently, the expected utility to an agent net of data collection costs is

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

Now assume there is a principal interested in acquiring an estimate at a specific precision denoted as θ_0 . She will choose a value for α so that the agent's precision is equal to θ_0 . That is, the principal selects an α which maximises the agent's expected utility at θ_0 by solving $\left. \frac{d\bar{U}}{d\theta} \right|_{\theta_0} = 0$ to give

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

The β parameter serves only to safeguard participation in the mechanism by ensuring that the agent's expected utility is positive. Presuming that the expected utility from the data collection and reporting alone shall be at least 0 we get

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

Based on Equations 2 and 3 we calculate the specific values of α and β which also depend on the strictly proper scoring rule used. This raises the important issue of which one of the most common rules (i.e., quadratic, spherical, logarithmic and parametric) the principal should select given that they have additional properties besides incentivising truthful reporting and eliciting sufficient effort. Papakonstantinou (2010) provides a thorough comparison of the four rules used as the basis of a second price single-dimensional auction including their analytical forms for Gaussian distributions and their bounds. It can be seen that the logarithmic rule and the parametric one with the parameter converging to 1, lead to the lowest expected payments, but they do not have lower bounds. This may result to un-realistically high penalties to agents that generate extremely inaccurate estimates with very low likelihood (i.e., $\mathcal{F} \rightarrow 0$). However, the parametric scoring rule offers a good compromise for values of the parameter within

(1, 1.5) as it can keep the payment relatively low for the majority of the agents, and the finite lower bound protects agents with extremely poor estimates.

4. The mechanism

The proposed mechanism implements a two-step payment to the agent that is selected to complete the procured project. That agent, is the one with the highest score based solely on his reported beliefs of his quality. In this context, the first payment to the selected agent is equal to the second score auction's payment based on his reported quality and the corresponding cost. Once that agent produces his quality and it is observed by the principal, he receives a secondary payment. This payment penalises inaccurate reports, and compensates for the data collection activities and the actual production.

The mechanism is formally defined as follows:

1. Principal invites N agents to participate in the procurement auction and announces that she needs estimates of their quality levels at precision θ_0 .
2. Agents generate estimates with means y_i and precisions θ_i , and report \hat{y}^i , $\hat{\theta}^i$ and the corresponding production costs \hat{x}^i with $i \in \{1, \dots, N\}$.
3. Each bid is assigned a score $Z^i = Z(\hat{x}^i, \hat{y}^i) = V(\hat{y}^i) - \hat{x}^i$.
4. The principal allocates the project to the agent with the highest score.
5. The selected agent receives his **first payment** from the principal: $P_A = P_A(\hat{x}^{(1)}, \hat{y}^{(1)}) = V(\hat{y}^{(1)}) - Z(\hat{x}^{(2)}, \hat{y}^{(2)})$ similar to the payment in a second score auction¹.
6. The selected agent produces quality y_0 .
7. The principal observes the agent's realised quality and issues the **second payment**:

$$P_B = P_B(y_0 | \hat{y}^{(1)}, \theta_0, \hat{x}^{(1)}) = d(y_0 | \hat{y}^{(1)}, \theta_0) P_A +$$

¹Superscripts (1) and (2) correspond to the agent with the highest and second highest score respectively.

$$+\alpha S(y_0|\hat{y}^{(1)}, \hat{\theta}^{(1)}) + \beta + V(y_0) - Z(\hat{x}^{(2)}, \hat{y}^{(2)})$$

where $d(y_0|\hat{y}^{(1)}\theta_0)$ is a function that evaluates the selected agent's reported estimate based on the observed actual production, cf. below, parameters α and β are the effort inducing parameters from the scaled strictly proper scoring rule $S(y_0|\hat{y}^{(1)}, \hat{\theta}^{(1)})$.

Like parameters α and β , the function $d(\cdot)$ serves to guarantee truthful reporting and that agents will generate estimates at the required precision in a setting where the agents' reports can deviate from their actual production due to unforeseen circumstances (endemic uncertainty) and due to strategic behaviour. In order to address both these effects, the deviation function $d(\cdot)$ is set equal to

$$d(y_0|\hat{y}^{(1)}, \hat{\theta}^{(1)}) = S(y_0|\hat{y}^{(1)}, \hat{\theta}^{(1)}) - \bar{S}(\theta_0) - 1 \quad (4)$$

where θ_0 is the precision required by the principal, $S(y_0|\hat{y}^{(1)}, \theta_0)$ is the scoring rule and $\bar{S}(\theta_0)$ is the expected score as a function of the required precision.

The *total* payment a truthful agent expects to derive by this mechanism is

$$\bar{P}(\theta) = [\bar{S}(\theta) - \bar{S}(\theta_0)][V(y) - Z^{(2)}] + \alpha \bar{S}(\theta) - \beta + V(y_0) - Z^{(2)} \quad (5)$$

In order to demonstrate how the mechanism works in practice, we provide a simple conceptual example while taking note that a more elaborate case study is presented in Section 6.1. Let us consider a business owner (principal) contracting the construction of a professional website. The firms participating in the procurement have received some specifications, while much of the technical details e.g., bandwidth, storage space, are left on them to decide based on their estimates of the website's traffic, with the principal providing some guidelines on how to estimate it. Let the specifications and corresponding uncertainty be aggregated in one quantity (referred as 'quality') and the guidelines on how to estimate traffic expressed by the required precision θ_0 . For the purpose of this example, we assume that the quality estimates are modelled as Gaussian distributions. Having generated their estimates and calculated the production costs, the firms report them with the principal receiving three

bids: $(9, 1, 72.9), (4, 1, 16), (16, 1, 224)$ in the form of (mean, precision, production cost) whereby production costs are calculated based on $x(y) = c_i * y^2$ with $c_i = \{0.9, 1, 0.875\}$.

Following the procedure of the mechanism, the principal calculates the scores and procures the service to the firm with the highest score. If the firms' value function is given by $V(y) = 50 * \sqrt{y} + 25$, the corresponding scores will be $Z^i = \{102.1, 109, 1\}$ (Steps 3 and 4). The firm with the highest score (firm 2) receives the contract and the first payment, P_A , equal to $125 - 102.1 = 22.9$ (Step 5) and then proceeds to create the contracted website, that is to produce quality equal to 4.2 (Step 6). Using the logarithmic scoring rule and Gaussian distributions to model the estimates, parameters α and β are calculated based on Papakonstantinou et al. (2011) as $2c$ and $c(2 - \ln \frac{1}{2\pi})$ respectively. To this end, assuming information collection cost for firm 2 roughly at the 1/100 of the production cost so that $C = 0.16$, it receives the second payment, P_B , equal to 14.73 (Step 7).

In the following section, we prove the mechanism's economic properties. After showing how the expression of the total payment in Equation 5 is derived, we prove that true revelation of an agent's estimates and costs is a Nash equilibrium and that the selected agent receives positive utility in expectation.

5. Economic properties

Having described in detail the mechanism, we now develop its economic properties. First, we show that the utility of a selected agent is maximised by generating an estimate at the principal's required precision (Lemma 1). Based on this result, we show that truthful revelation of the production costs is an optimal strategy, given the truthful report of the quality estimate generated by a sequence of consistent estimators (Lemma 2). Generalising these two specific results allows to show the mechanism's immunity to combined misreporting and that truth revelation is in-fact a Nash equilibrium. Specifically we show that:

1. Agents generate their estimates at the principal's required precision.

2. The mechanism is immune to the effects of combined misreporting of quality and cost with truth revelation being a Nash equilibrium.
3. The mechanism is individually rational for the selected agent.

Lemma 1. *The expected utility of the selected agent is maximised only by generating an estimate at the principal's required precision and truthfully revealing it.*

Proof. Given the mechanism described above, when an agent reports his estimate, he must do so with a precision up to the one required by the principal. The selected² agent's utility when he reports $(\hat{y}, \hat{\theta})$, given realised quality equal to y_0 , is

$$U(y_0 | \hat{y}, \hat{\theta}) = V(\hat{y}) - Z^{(2)} + [S(y_0 | \hat{y}, \hat{\theta}) - \bar{S}(\theta_0) - 1][V(\hat{y}) - Z^{(2)}] + \\ + \alpha S(y_0 | \hat{y}, \hat{\theta}) + \beta + V(y_0) - Z^{(2)} - x(y_0) - c(\theta)$$

where α and β are the strictly proper scoring rules scaling parameters defined in Section 3.

By integrating over the set of possible qualities y_0 we derive the expected utility from reporting $(\hat{y}, \hat{\theta})$:

$$\bar{U}(\hat{y}, \hat{\theta}) = \int_{-\infty}^{\infty} \mathcal{F}(y_0 | y, \theta) [V(\hat{y}) - Z^{(2)}] dy_0 \\ + \int_{-\infty}^{\infty} \mathcal{F}(y_0 | y, \theta) [S(y_0 | \hat{y}, \hat{\theta}) - \bar{S}(\theta_0) - 1][V(\hat{y}) - Z^{(2)}] dy_0 \\ + \int_{-\infty}^{\infty} \mathcal{F}(y_0 | y, \theta) [\alpha S(y_0 | \hat{y}, \hat{\theta}) + \beta] dy_0 \\ + \int_{-\infty}^{\infty} \mathcal{F}(y_0 | y, \theta) [V(y_0) - Z^{(2)} - x(y_0) - c(\theta)] dy_0$$

Since the initial payment does not depend on the final outcome and $\int_{-\infty}^{\infty} \mathcal{F}(y_0 | y, \theta) dy_0 = 1$, a shorter expression is derived:

$$\bar{U}(\hat{y}, \hat{\theta}) = [V(\hat{y}) - Z^{(2)}] \int_{-\infty}^{\infty} \mathcal{F}(y_0 | y, \theta) [S(y_0 | \hat{y}, \hat{\theta}) - \bar{S}(\theta_0)] dy_0$$

²In order to simplify our notation we omit the use of superscript (1) to denote the agent with the highest score, while we maintain the use of (2) for the runner-up agent. We will also be using the simplified version of the auction scores i.e. $Z^i = Z(\hat{x}^i, \hat{y}^i)$

$$+ \int_{-\infty}^{\infty} \alpha \mathcal{F}(y_0|y, \theta) S(y_0|\hat{y}, \hat{\theta}) dy_0 + \beta - c(\theta) + \int_{-\infty}^{\infty} \mathcal{F}(y_0|y, \theta) [V(y_0) - x(y_0)] dy_0 + Z^{(2)}$$

The above expression can be further simplified by using the notation of the expected score:

$$\bar{S}(\hat{\mathcal{F}}, \mathcal{F}) = \int_{-\infty}^{\infty} \mathcal{F}(y_0|y, \theta) S(y_0|\hat{y}, \hat{\theta}) dy_0$$

where $\hat{\mathcal{F}}$ represents the prediction corresponding to the reported parameters $(\hat{y}, \hat{\theta})$ and \mathcal{F} the actual distribution of the production quality.

To sum up, the selected agent's expected utility from estimating and reporting his quality is:

$$\begin{aligned} \bar{U}(\hat{y}, \hat{\theta}) &= [V(\hat{y}) - Z^{(2)}][\bar{S}(\hat{\mathcal{F}}, \mathcal{F}) - \bar{S}(\theta_0)] + \alpha \bar{S}(\hat{\mathcal{F}}, \mathcal{F}) + \beta \\ &+ \int_{-\infty}^{\infty} \mathcal{F}(y_0|y, 1/\theta) [V(y_0) - x(y_0)] dy_0 + Z^{(2)} - c(\theta) \end{aligned}$$

Having defined the selected agent's expected utility function we proceed to show that it is maximised when the agent's reported estimate is equal to his true one, generated at a precision equal to the principal's required one. Initially, it is easy to see that due to the use of a strictly proper scoring rule, the expected scoring rule $\bar{S}(\hat{\mathcal{F}}, \mathcal{F})$ is maximised at $\hat{\mathcal{F}} = \mathcal{F}$, hence (y, θ) is a local maximum for the expected score.

In order to show that (y, θ) is a maximum point for \bar{U} , we first show that it is a critical one,

$$\frac{\partial \bar{U}}{\partial \hat{y}} = V'(\hat{y})[\bar{S}(\hat{\mathcal{F}}, \mathcal{F}) - \bar{S}(\hat{\theta})] + [V(\hat{y}) - Z^{(2)} + \alpha] \frac{\partial \bar{S}(\hat{\mathcal{F}}, \mathcal{F})}{\partial \hat{y}} = 0 \quad (6)$$

$$\frac{\partial \bar{U}}{\partial \hat{\theta}} = [V(\hat{y}) - Z^{(2)} + \alpha] \frac{\partial \bar{S}(\hat{\mathcal{F}}, \mathcal{F})}{\partial \hat{\theta}} = 0 \quad (7)$$

From the definition of a strictly proper scoring rule, the partial derivatives of $\bar{S}(\hat{\mathcal{F}}, \mathcal{F})$ w.r.t \hat{y} and $\hat{\theta}$ are equal to 0, for $(\hat{y}, \hat{\theta}) = (y, \theta)$, hence the first order conditions for \bar{U} are fulfilled by (y, θ) .

Moreover, the determinant of the Hessian matrix of \bar{U} is

$$\text{Det}(\mathcal{H}(\bar{U}))(y, \theta) = [V(y) - Z^{(2)} + \alpha]^2 \left[\frac{\partial^2 \bar{S}}{\partial \hat{y}^2} \frac{\partial^2 \bar{S}}{\partial \hat{\theta}^2} - \frac{\partial^2 \bar{S}}{\partial \hat{y} \partial \hat{\theta}} \frac{\partial^2 \bar{S}}{\partial \hat{\theta} \partial \hat{y}} \right] =$$

$$= [V(y) - Z^{(2)} + \alpha]^2 \text{Det}(\mathcal{H}(\bar{S}))(y, \theta) \quad (8)$$

which is positive given that $[V(y) - Z^{(2)} + \alpha]^2 > 0$ and $\text{Det}(\mathcal{H}(\bar{S}))(y, \theta) > 0$ since (y, θ) is a maximum for the expected score \bar{S} .

Having shown that (y, θ) is a maximum for the selected agent's expected utility, it is trivial to prove based on the definition of the parameters α and β in Section 3 that a selected agent can maximise his utility by making an estimate at a certain precision, in this particular case equal to principal's requirement of θ_0 . In such case, the utility he expects to derive is equal to

$$\bar{U}(y) = \int_{-\infty}^{\infty} \mathcal{F}(y_0|y, 1/\theta)[V(y_0) - x(y_0)]dy_0 - Z^{(2)} \quad (9)$$

□

Lemma 2. *The use of consistent estimators for the respective outcomes, y_0^i , guarantees that agents' expected utilities are maximised by truthfully revealing their production cost given that they have truthfully revealed their quality estimates.*

Proof. As already mentioned in Section 2, we model uncertainty in agents' qualities by assuming individual estimates as parametric distributions with means y^i and precision θ^i such that $y^i \sim \mathcal{F}(y_0^i, \theta_i)$. Now, the use of a consistent sequence of estimators for the estimation of y_i^0 e.g. a sequence of sample means (DeGroot & Schervish, 2002) for the estimation of the mean of a Gaussian distribution (i.e. $\mathcal{F} = \mathcal{N}$)³ suggest that for a sufficiently large sample of observations the utility a truthful selected agent expects to derive is equal to:

$$\bar{U}(y) = \int_{-\infty}^{\infty} \mathcal{F}(y_0|y, 1/\theta)[V(y_0) - x(y_0)]dy_0 - Z^{(2)} = V(y) - x(y) - Z^{(2)} = Z^{(1)} - Z^{(2)}$$

Based on the use of a consistent estimator, we proceed to prove this Lemma by contradiction:

³In Section 6 we show that this assumption does not limit the implementation of the mechanism based on an example with Gaussian distributions and a Bayesian estimator of its mean. We evaluate the estimator's convergence and its effect on the mechanism.

Let x and y be an agent's true cost and quality, Z the score that corresponds to these true values, and let \hat{x}, \hat{y} be the reports and \hat{Z} the score corresponding to these reports. Furthermore, let $x^{(2)}, y^{(2)}, Z^{(2)}$ be the bids, and the score of the runner up agent (i.e. $\hat{Z} > Z^{(2)}$).

First, let an agent's misreporting of his cost have an impact on the outcome of the auction. We consider the following two cases:

1. Agent wins by misreporting while it would have lost if truthful.
 2. Agent loses by misreporting while it would have won if truthful.
- In Case (1) agent reports his cost s.t $\hat{Z} > Z^{(2)}$ given that $Z < Z^{(2)}$. The agent achieves this by reporting a lower cost than his actual one i.e. $\hat{x} < x$. Under optimal reporting of quality, the utility of an agent misreporting his cost in Case (1) will be negative i.e. $\bar{U}(y) = V(y) - x(y) - Z^{(2)} = Z^{(1)} - Z^{(2)} < 0$.
 - In Case (2) agent reports his cost such that $\hat{Z} < Z^{(2)}$ given that $Z > Z^{(2)}$. The agent would have won the auction, and therefore would be selected, but instead reports a cost greater than his actual one i.e. $\hat{x} > x$. As a result, the agent loses the auction and consequently receives negative utility (since he still faces the costs of determining his quality).

Second, we assume that the agent misreports his cost of production without this affecting whether he wins the auction or not. If the agent had already lost the auction, misreporting would have no additional effect given that the utility would be negative due to the cost of determining his quality without any dependence on the cost of production. Had the agent already won the auction, misreporting would not result in additional benefits. Specifically, both payments depend on the second lowest score, that is, the payment for the reported and actually produced (for the second stage) quality and the compensation for his estimate. \square

Theorem 1. *The mechanism is immune to combined misreporting of quality and cost.*

Proof. In the above proofs we showed that a selected agent maximises his expected utility by truthfully reporting his production quality if he reports truthfully his cost, and that the same holds for an agent's costs, given that he generated an accurate estimate of his quality by investing the required amount of resources in determining it. However, given the multi-dimensional nature of the bids an agent could attempt to manipulate the principal by misreporting both costs and the precision of his quality estimate.

In this proof we examine agents' combined strategic behaviour. In order to demonstrate how it is not optimal for an agent to deviate from truthful behaviour we consider the four following general cases of misreporting:

1. Agent wins the auction by misreporting both his estimate of quality and production cost
2. Agent wins the auction with the misreporting having no effect on the auction's outcome
3. Agent loses the auction due to his misreporting
4. Agent loses the auction despite his misreporting

- In Case (1) the agent reports his estimate of quality and cost such that $\hat{Z} > Z^{(2)}$, while $Z < Z^{(2)}$, with the precision of the estimate θ not necessarily equal to the principal's θ_0 . We will show that the misreporting agent's expected utility $\bar{U}(\hat{y}, \hat{\theta})$ will always be less or equal to the utility of a truthful agent $\bar{U}(y, \theta_0)$ following the principal's instructions:

$$\bar{U}(\hat{y}, \hat{\theta}) - \bar{U}(y, \theta_0) = [V(\hat{y}) - Z^{(2)}][\bar{S}(\hat{\mathcal{F}}, \mathcal{F}) - \bar{S}(\theta_0)] + \alpha \bar{S}(\hat{\mathcal{F}}, \mathcal{F}) + \beta - c(\theta) \quad (10)$$

Regarding $V(\hat{y}) - Z^{(2)}$ we have assumed that it is a positive quantity since $\hat{Z} > Z^{(2)} \Rightarrow V(\hat{y}) - \hat{x}(\hat{y}) > Z^{(2)} \Rightarrow V(\hat{y}) > Z^{(2)}$, while $\bar{S}(\hat{\mathcal{F}}, \mathcal{F}) - \bar{S}(\theta_0)$ is negative since $\bar{S}(\hat{\mathcal{F}}, \mathcal{F}) \leq \bar{S}(\theta_0)$ given that S is a strictly proper scoring rule. Finally, from Section 3 it can be trivially shown that $\alpha \bar{S}(\hat{\mathcal{F}}, \mathcal{F}) + \beta - c(\theta) < 0$ for $\theta < \theta_0$; hence $\bar{U}(\hat{y}, \hat{\theta}) \leq \bar{U}(y, \theta_0)$.

- In Case (2) the agent would have won the auction anyway, and although misreporting of cost and quality will have no impact on the outcome of the

auction, it may have on the secondary payment. Still, such manipulation is not attractive since we have from Case (1) that $\bar{U}(\hat{y}, \hat{\theta}) \leq \bar{U}(y, \theta_0)$. Even if we assume that the estimate's precision is equal to required one by the principal, it is still the misreporting of the estimate and the production cost which makes this strategy sub-optimal.

- Cases (3) and (4) are simpler. For both cases it is obvious that the utility of an agent not winning the auction will solely consist of the cost of data collection. Specifically, in Case (3) the agent deliberately misreports his estimate and his production cost in order to lose. It would be in his best interest to invest minimum resources in generating his estimate, so that he can minimise his inevitable loss. However, that is not a straightforward decision. Estimates of low precision may end up winning the auction and inflicting additional losses, while estimates of high precision will increase his losses. Effectively, an agent that wants to lose the auction has no reason to participate in it, specially given the assumption of rational and risk neutral participants.

Now, in Case (4) the agent misreports with the intention to win but ends up losing the auction. Had the agent won, it would result in negative utility as shown in Case (1) and given that the agent intends to win, he will invest the required resources in generating his estimate, hence his loss due to prediction costs will increase.

Having shown that combined misreporting of costs, estimates of qualities and their precision leads to sub-optimal outcomes including negative utilities, we proved that the mechanism is immune to this type of strategic behaviour. In conjunction with Lemma 2, this leads to the following corollary: \square

Corollary 1. *Truth revelation of an agent's quality and cost is a Nash equilibrium.*

Proof. A selected agent maximises his expected utility by reporting his true production costs and qualities which are generated at the precision required

by the principal. Nevertheless, Lemma 2 imposes an additional constraint by proving that an agent will reveal his true production cost under the assumption that he has truthfully revealed his quality estimate. For this to hold for an individual agent it must extend to all others, thus weakening the equilibrium of the proposed mechanism and moving it from dominant strategy to Nash. \square

Theorem 2. *The use of consistent estimators for the respective outcomes, y_0^i , guarantees that the utility of the selected agent is positive in expectation.*

Proof. The utility an agent that has truthfully reported his estimates, his precisions and the productions costs, expects to derive is given by

$$\bar{U}(y) = V(y) - x(y) - Z^{(2)} = Z^{(1)} - Z^{(2)}$$

Given that $V(y) - x(y)$ is the selected agent's true score, the utility is positive on expectation, hence the mechanism is interim individually rational for that specific agent. \square

It should be noted that the above proof holds for any strictly proper scoring rule. However, it is possible to move to a stronger individual rationality concept (ex-post individual rationality) by restricting the analysis to the four most well known strictly proper scoring rules (i.e. quadratic, spherical, logarithmic and parametric). This leads to the following corollary:

Corollary 2. *The use of consistent estimators for the respective outcomes, y_0^i , guarantees that the utility of the selected agent is positive for each one of the quadratic, spherical, logarithmic and parametric scoring rules.*

Proof. The use of a consistent sequence of estimators implies that for relative large samples of observations during the estimation process the selected agent's estimate will be very close to the actual outcome y_0 , effectively allowing to replace y with y_0 . After this transformation, the selected agent's utility function takes the following form

$$U(y_0) = [S(y_0) - \bar{S}(\theta) + 1][V(y_0) - Z^{(2)}] + \alpha S(y_0) + \beta - x(y_0) - c(\theta) - x(y_0) =$$

$$[S(y_0) - \bar{S}(\theta)][V(y_0) - Z^{(2)}] + [S(y_0) - \bar{S}(\theta)] \frac{c'(\theta)}{\bar{S}'(\theta)} + Z^{(1)} - Z^{(2)}$$

We have already shown that $V(y_0) - Z^{(2)} > 0$ and $Z^{(1)} - Z^{(2)} > 0$; also given that both $c(\cdot)$ and $\bar{S}(\cdot)$ are increasing $\frac{c'(\theta)}{\bar{S}'(\theta)} > 0$. Based on the above and on the analytic calculations of $S(y_0)$ for the four strictly proper scoring rules listed in Section 3 and of the expected values $\bar{S}'(\theta)$ derived by Papakonstantinou et al. (2011) it can be seen that the expression of the selected agent's utility is indeed positive. \square

6. Numerical evaluation

In this section we evaluate the robustness of the mechanism in a scenario which does not involve the theoretical assumptions regarding a sufficiently large sample of observations, since this may not be plausible in an applied setting. In order to do so, first we explain the sampling and estimation process and then proceed to demonstrate how this mechanism works through an illustrative example. We then undertake a series of simulations to get a better understanding of its performance. We compare our mechanism with two benchmark cases and highlight the costs of uncertainty regarding the agents' predictions of their quality.

Specifically, as the first benchmark case, labelled as 'Second Score: Outcome', we use the standard second score auction (SSA) under the assumption that there is no uncertainty and agents can directly report their actual qualities y_0^i . For the second case, labelled as 'Second Score: Belief', we introduce uncertainty in the model and modify the second score auction so it uses the agents' beliefs of their qualities instead of the actual outcomes of the previous case, under the assumption that the principal can elicit their truthful reporting without external means such as fixed penalties and fines.

Against this background, let agent i attempt to estimate his individual quality y_0^i by generating a sample of M independent observations $y_j^i, j \in \{1, \dots, M\}$. As already mentioned, agents do not know their quality ex ante; instead they

have an a-priori belief which they update by collecting additional information. We assume that agent i 's a-priori belief about y_0^i is given as $y_0^i \sim \mathcal{N}(y_\mu^i, 1/\theta_\mu^i)$, and that he is able to collect further information about y_0^i by generating M independent and identically distributed random observations $\{y_1^i, y_2^i, \dots, y_M^i\}$ with $y_j^i \sim \mathcal{N}(y_0^i, 1/\theta_j^i)$. Using these observations, the agent can update his a-priori beliefs to a posterior belief

$$y_0^i \sim \mathcal{N}\left(\frac{y_\mu^i \theta_\mu^i + \theta^i \bar{y}^i}{\theta_\mu^i + \theta^i}, \frac{1}{\theta_\mu^i + \theta^i}\right) \quad (11)$$

where \bar{y}^i is the mean of the observations $\{y_1^i, y_2^i, \dots, y_M^i\}$ and θ^i is the resulting precision of the sample average \bar{y}^i , i.e. $\theta^i = M\theta_j^i$. By sampling with precision θ^i , and spending information collection costs $c^i(\theta^i)$ it is able to predict its quality y^i as $(y_\mu^i \theta_\mu^i + \bar{y}^i \theta^i)/(\theta_\mu^i + \theta^i)$ with precision $\theta_\mu^i + \theta^i$.

In the simulations, we consider a specific case in which the parameters y_μ^i of the agents' prior beliefs of their production qualities are drawn from the uniform distribution $\mathcal{U}(2, 5)$, while we assume that the agents' precisions in both priors and individual observations during data collection are equal to 1. Consequently, given the parameters of our model, the actual production quality level follows the Gaussian distribution $\mathcal{N}(y_\mu^i, 1)$. Furthermore, the agents' production cost functions are given by $x^i = X^i y^2$, where $X^i \sim \mathcal{U}(0, 1)$, while the costs of data collection are linear functions, given by $c^i = C^i \theta$, where $C^i \sim \mathcal{U}(0.001, 0.002)$. Note that the bounds in the distribution of the data collection cost parameter are selected so that even for relative large samples the overall cost is relative small compared to the actual production cost. A scenario whereby data collection cost would be higher than the production cost is not considered to be neither interesting nor realistic.

The principal's value function is given by $V(y) = B(1 - e^{-y})$, with $B = 20$ guaranteeing that there will be some agents with positive scores $V(y) - x(y)$ in the range of qualities we use. This particular value function is both increasing and concave and it provides some curvature, as opposed to a conventional approach such as $V(y) = B\sqrt{y}$ which is almost linear in the range of B that

guarantees positive scores.

The mechanism is simulated 10^5 times, with 20 agents participating in the auction, while the precision of each agent’s sample average, and consequently his sample of observations, M , ranges from 1 to 100. Our theoretical analysis suggests that the principal’s required precision θ_0 also takes values from 1 to 100. For each iteration we record the selected agent’s utility, his payment, his prediction and production costs and whether the agent selected by our mechanism is the agent that would have been selected had there been no uncertainty (we refer that agent as a ‘proper winner’). For the calculations that involve a lack of uncertainty, the agents will report their actual outcome y_0^i directly. In a given iteration all agents face underlying cost functions of the same form, but their priors, sample observations and cost parameters differ. Due to the number of iterations the standard error in the mean values plotted is in the range of 10^{-4} to 10^{-5} and thus we omit the use of error bars for clarity.

6.1. An illustrative example

For a snapshot of the mechanism, we calculate several of the mechanism’s elements i.e. selected agents, payments and costs as the sample’s precision increases. Specifically, Table 1 lists the selected agent in our mechanism and the winner of the second score auction with no uncertainty (i.e. Second Score: Outcome), denoted as w and w' respectively. We also calculate the parts of the secondary payment i.e. the d function: $S(y_0|\hat{y}, \hat{\theta}) - \bar{S}(\theta_0) - 1$ and the penalty for inaccuracies: $d(y_0|\hat{y}, \hat{\theta})[V(\hat{y}) - Z^{(2)}]$, while listing the first and secondary payments of the mechanism (Steps 6 and 8 respectively), the total payment and the winner’s utility. Finally, the last column is the ratio between the cost of production $x(y)$ and data collection $c(\theta_0)$.

From Table 1 it can be seen that in this particular instance, and at a required sample precision of 4, the selected agent from our mechanism is the winner of the second score auction with no uncertainty, $w = w'$. This shows that our mechanism identified the ‘proper’ winner, i.e. the agent who should have won based solely on actual production, after he generated a sample of 4 observations.

Table 1: A single iteration of the mechanism.

θ_0	w	w'	$d(\cdot)$	Penalty	1st Pay	2nd Pay	Total P	Utility	Cost ratio
1	11	14	-1.5862	-1.7624	1.1111	-3.4723	-2.3612	-2.5621	110.09
4	14	14	-0.9166	-0.5276	0.5756	-0.1340	0.4417	0.0692	49.88
10	14	14	-0.9125	-0.9636	1.0560	-0.0171	1.0389	0.6554	19.95
16	14	14	-0.8846	-0.7820	0.8840	0.1577	1.0417	0.6472	12.47
22	14	14	-0.9417	-0.6751	0.7169	0.2874	1.0043	0.5989	9.07
28	14	14	-0.8786	-0.7433	0.8460	0.2266	1.0726	0.6562	7.13
34	14	14	-0.9168	-0.9277	1.0119	0.0906	1.1024	0.6751	5.87
40	14	14	-0.9382	-0.9601	1.0233	0.0581	1.0814	0.6430	4.99
46	14	14	-0.9213	-0.9190	0.9975	0.1251	1.1225	0.6732	4.34
52	14	14	-0.8727	-0.8030	0.9201	0.3023	1.2224	0.7621	3.84
58	14	14	-0.9004	-0.7504	0.8334	0.3573	1.1907	0.7194	3.44
64	14	14	-0.8758	-0.7663	0.8750	0.3885	1.2635	0.7812	3.12
70	14	14	-0.8820	-0.7550	0.8561	0.4070	1.2631	0.7698	2.85
76	14	14	-0.8658	-0.7755	0.8958	0.4239	1.3197	0.8154	2.63
82	14	14	-0.8671	-0.7665	0.8840	0.4551	1.3391	0.8239	2.43
88	14	14	-0.8638	-0.7765	0.8990	0.4678	1.3668	0.8406	2.27
94	14	14	-0.8645	-0.7864	0.9096	0.4791	1.3887	0.8515	2.12
100	14	14	-0.8687	-0.7578	0.8724	0.5239	1.3964	0.8482	2.00

However, the $d(\cdot)$ function is not equal to -1 , as it is in expectation, which leads to a heavier fine for the selected agent of our mechanism. Hence the 2nd Pay, total payment and utility are negative for some values of θ_0 . Specifically regarding the selected agent's utility, it is interesting to observe the loss of an imprecise agent, and the relation with our theoretical results in Section 5, where we discussed how estimates of low precision may end up winning the auction but inflicting additional losses instead of gains (Theorem 1, Case (3)). Still, despite the good intuition that this analysis provides for our mechanism, it should be

noted that these results are from a single iteration, hence exposed to heavy bias from the random inputs (i.e. costs and qualities).

6.2. Numerical simulations




Figure 2: Dependence of percentage of ‘proper’ winners and expected utility on sample precision.

Having detailed the simulation’s input parameters and analysed a snapshot of the mechanism, we now present our numerical findings after simulating the mechanism 10^5 times. In Fig. 2 we summarise the convergence of our mechanism. It can be seen that for the specific scenario we consider, it takes a relatively small required precision, for the outcome of our mechanism to be the same as the outcome of the second score auction with no uncertainty, where agents directly report their realised qualities (i.e. Second Score: Outcome).

In fact, for required precisions above the area of 50 the selected agent from our mechanism and the winner of the second score auction are the same agents in more than 95% of the iterations of the mechanism (Fig. 2(a)). In addition to this, our analytic findings in Section 5 are validated in Fig. 2(b), where we notice that selected agent’s expected utility increases as the required precision increases. The utility the selected agent expects to derive is less than the second score auction’s winner expected utility (Second Score: Belief), had he been able

to generate and report his belief of his quality freely. As it is expected, as the required precision increases both auctions approach the second score auction in a setting with no uncertainty where the winner can report his actual production from the beginning (Second Score: Outcome). The differences that appear are attributed to those cases where the winners of the two auctions do not coincide, hence the winner faces losses.




Figure 3: The selected agent’s expected payment for both production and prediction and ratio of production to prediction costs.

The payment the selected agent in our mechanism and in the conventional benchmark expects to derive, as well as, his average costs for the required precision $\theta_0 \in [1, 100]$ are shown in Fig. 3. The penalties for inaccuracies, and the principal’s compensation for the data collection costs have a clear impact on the expected payment as shown in Fig. 3(a). The expected payment in our mechanism starts lower than the two benchmark auctions, but it increases as the precision increases. The stability in the payments of the other auctions is to be expected since there is no data collection before the auction hence no compensation, while the higher payments for our mechanism will not be an issue in realistic applications since the data collection costs tend to be significantly lower than the production costs. Nevertheless, it should be noted that the payments’ differences are highlighted in the plot due to its scale, with this issue related to

the particular selection of simulation parameters and not to the mechanism itself. In fact, even after setting the upper bound of the cost collection parameter equal to 0.002 and using a linear cost function, for relatively high precisions, information collection cost ends up very close to some agents' production costs. We demonstrate this data sensitivity in Fig. 3(b), where we plot the logarithmic ratio of production to prediction costs.

7. Conclusions

In this article we introduced a novel mechanism for multi-dimensional procurement auctions where the agents' qualities are uncertain and the principal can only verify them after the project is completed. The principal cannot enforce truthful reporting of costs and qualities through external means nor she can force agents to provide quality-cost estimates of extremely high precisions.

Initially the principal procures an item or a service from a group of agents by implementing a standard second score auction, only now the agents' ranking is calculated based on their reported estimates of their qualities and costs. The agent with the higher score is allocated the contract to provide the agreed item or service and receives the second score payment. After he fulfils his part of the contract he receives a secondary payment based on both the reported estimate and the actual production observed by the principal. The secondary payment uses a strictly proper scoring rule to evaluate the selected agent's posterior belief of his quality once the project is complete and the principal can witness the outcome.

We showed that the mechanism is immune to the agents' combined misreporting of quality and costs. In addition to that, we showed that participating agents invest the amount of the resources required when generating the estimates, while they expect to have a positive utility if they are allocated the project. Although, the rest of the agents incur losses due to the information collections costs, the payment to the agent that receives and completes the project at his reported quality exceeds the production and information costs, provided

that the estimate is of relative high precision.

As the assumption about the required precision indicates, there are some limitations regarding practical elements of the mechanism. Although we proved analytically that our mechanism implements the standard second score auction's outcome in terms of the selected agent's expected utility, numerical simulations demonstrated how sensitive the mechanism is to the precision of the quality predictions, and hence to the resources invested in preparing the bids. This calls for extra attention by the principal; the required precision must be set at a level whereby the theoretical properties of the mechanism are not compromised while the agents are willing to participate and expect to cover more than their data collection and production cost if they are allocated the project. Therefore, an early 'consultation' stage may be required to set the required precision based on prior history of industry practices.

A final practical issue of the proposed mechanism is its potential vulnerability to collusion; a possibility, even if such behaviour is often considered illegal by regulating authorities. Literature has identified several cases where single-dimensional auctions, which in turn can be extended to their multi-dimensional counterparts introduced by Che (1993), can be vulnerable to collusion. However, even for single-dimensional auctions collusion is not a straightforward process. In more detail, Mailath & Zemsky (1991) analyse the non-trivial process of identifying an auction winner and calculating payments to the losers within a cartel. More recently, Che & Kim (2009) introduced third-party uninformed representatives to coordinate the formation and actions of multiple cartels, without really considering practicalities such as how the actual side-payments will be issued so that that they do not raise suspicion to regulation authorities.

The fact that efficient collusion assumes some form of coordination within the cartels suggests that our mechanism could be less vulnerable mainly due the additional complexity brought by the uncertain qualities. Bargaining within the cartel now is an even more complex process given that agents have to invest costly resources in estimating their qualities and then be held accountable for that uncertainty not only by the principal, but also by the other participants in

their cartel. Naturally, bargaining under uncertain and asymmetric information can be less efficient than bargaining under perfect information. Such complex behaviour exceeds the scope of this paper and provides a very interesting and particularly challenging line of future work.

References

- Anderson, S. P., Goeree, J. K., & Holt, C. A. (1998). Rent seeking with bounded rationality: An analysis of the allpay auction. *Journal of Political Economy*, 106, 828–853.
- Asker, J., & Cantillon, E. (2008). Properties of scoring auctions. *RAND Journal of Economics*, 39, 69–85.
- Beil, D. R., & Wein, L. (2003). An inverse-optimization-based auction mechanism to support a multiattribute rfq process. *Management Science*, 49, 1529–1545.
- Bichler, M. (2000). An experimental analysis of multi-attribute auctions. *Decision Support Systems*, 29, 249–268.
- Bichler, M., & Kalagnanam, J. (2003). Configurable offers and winner determination in multi-attribute auctions. *European Journal of Operational Research*, 160, 380–394.
- Bogetoft, P., & Nielsen, K. (2008). Dea based auctions. *European Journal of Operational Research*, 184, 685–700.
- Branco, F. (2007). The design of multidimensional auctions. *RAND Journal of Economics*, 28, 63–81.
- Charnes, A., Cooper, W. W., & Rhodes, E. (1979). Short communication: Measuring the efficiency of decision making units. *European Journal of Operational Research*, 3, 339.

- Charnes, A., Cooper, W. W., & Rhodes, E. (1987). Measuring the efficiency of decision making units. *European Journal of Operational Research*, 2, 429–444.
- Che, Y.-K. (1993). Design competition through multidimensional auctions. *RAND Journal of Economics*, 24, 668–680.
- Che, Y.-K., & Kim, J. (2009). Optimal collusion-proof auctions. *Journal of Economic Theory*, 144, 565–603.
- Clarke, E. (1971). Multipart pricing of public goods. *Public Choice*, 11, 17–33.
- DeGroot, M. H., & Schervish, M. J. (2002). *Probability and Statistics*. Addison Wesley.
- Friedman, D. (1983). Effective scoring rules for probabilistic forecasts. *Management Science*, 29, 447–454.
- Gregory, P. C. (2005). *Bayesian Logical Data Analysis for the Physical Sciences: A Comparative Approach with Mathematica Support*. Cambridge University Press.
- Groves, T. (1973). Incentives in teams. *Econometrica*, 41, 617–631.
- He, M., Jennings, N. R., & Leung, H.-F. (2003). On agent-mediated electronic commerce. *IEEE Transactions on knowledge and data engineering*, 15, 985–1002.
- Hendrickson, A. D., & Buehler, R. J. (1971). Proper scores for probability forecasters. *The Annals of Mathematical Statistics*, 42, 1916–1921.
- Jurca, R., & Faltings, B. (2005). Reputation-based service level agreements for web services. In *Service Oriented Computing* (pp. 396–409). Springer Berlin / Heidelberg volume 3826 of *Lecture Notes in Computer Science*.
- Jurca, R., & Faltings, B. (2006). Minimum payments that reward honest reputation feedback. In *Proceedings of the ACM Conference on Electronic Commerce* (pp. 190–199). Ann Arbor, Michigan, USA.

- Jurca, R., & Faltings, B. (2007). Collusion resistant, incentive compatible feedback payments. In *Proceedings of the ACM Conference on Electronic Commerce* (pp. 200–209). San Diego, California, USA.
- Liu, T. X., Yang, J., Adamic, L. A., & Chen, Y. (2014). Crowdsourcing with all-pay auctions: A field experiment on taskcn. *Management Science*, 60. doi:10.1287/mnsc.2013.1845.
- Mailath, G. J., & Zemsky, P. (1991). Collusion in second price auctions with heterogeneous bidders. *Games and Economic Behavior*, 3, 467–486.
- McAdams, D. (2006). Monotone equilibrium in multi-unit auctions. *Review of Economic Studies*, 73, 1039–1056.
- Miller, N. H., Pratt, J. W., Zeckhauser, R. J., & Johnson, S. (2007). Mechanism design with multidimensional, continuous types and interdependent valuations. *Journal of Economic Theory*, 136, 476–496.
- Miller, N. H., Resnick, P., & Zeckhauser, R. J. (2005). Eliciting honest feedback: The peer prediction method. *Management Science*, 51, 1359–1373.
- Papakonstantinou, A. (2010). *Mechanism Design for Eliciting Costly Observations in Next Generation Citizen Sensor Networks*. Ph.D. thesis University of Southampton, School of Electronic and Computer Science.
- Papakonstantinou, A., & Bogetoft, P. (2013). Incentives in multi-dimensional auctions under information asymmetry for costs and qualities. In *Agent-Mediated Electronic Commerce. Designing Trading Strategies and Mechanisms for Electronic Markets - AMEC / TADA 2012* (pp. 104–118). Springer volume 136 of *Lecture Notes in Business Information Processing*.
- Papakonstantinou, A., Rogers, A., Gerding, E. H., & Jennings, N. R. (2011). Mechanism design for the truthful elicitation of costly probabilistic estimates in distributed information systems. *Artificial Intelligence*, 175, 648–672.

- Parkes, D., & Kalagnanam, J. (2005). Models for iterative multiattribute procurement auctions. *Management Science*, *51*, 435–451.
- Savage, L. J. (1977). Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association*, *66*, 783–801.
- Vickrey, W. (1961). Counterspeculation, auctions and competitive sealed tenders. *The Journal of Finance*, *16*, 8–37.
- Witkowski, J., & Parkes, D. C. (2012). Peer prediction without a common prior. In *Proceedings of the 13th ACM Conference on Electronic Commerce EC '12* (pp. 964–981).