



Munich Personal RePEc Archive

Multi-dimensional auctions under information asymmetry for costs and qualities

Papakonstantinou, A. and Bogetoft, P.

Copenhagen Business School, Department of Economics

2012

Online at <https://mpra.ub.uni-muenchen.de/43563/>

MPRA Paper No. 43563, posted 04 Jan 2013 16:48 UTC

Incentives in Multi-dimensional Auctions under Information Asymmetry for Costs and Qualities¹

A. Papakonstantinou, P. Bogetoft

Economics Department, Copenhagen Business School
{ap.eco, pb.eco}@cbs.dk

Summary. This paper discusses the design of a novel multi-dimensional mechanism which allows a principal to procure a single project or an item from multiple suppliers through a two-step payment. The suppliers are capable of producing different qualities at costs which cannot exceed a certain value and the mechanism balances between the costs faced by the suppliers and the benefit the principal achieves from higher qualities. Initially, the principal implements a standard second score auction and allocates the project to a single supplier based its reported cost and quality, while then it elicits truthful reporting of the quality by issuing a symmetric secondary payment after observing the winner's production. We then provide an alternate mechanism in which the principal issues an asymmetric secondary payment which rewards agents for producing higher qualities, while it penalises them for producing lower qualities than they reported. We prove that for both mechanisms truthful revelation of costs and qualities is a dominant strategy (weakly for costs) and that they are immune to combined misreporting of both qualities and costs. We also show that the mechanisms are individually rational, and that the optimal payments received by the winners of the auctions are equal to the payment issued by the standard second score auction.

1.1 Introduction

The appearance and effectively the dominance of the Internet in many everyday activities has signaled a digital revolution which has thrown some considerable attention to auction mechanisms. Auctions are now playing an instrumental role in a vast number of on-line activities, whether affecting them directly (e.g. on-line trading sites such as ebay and taobao) or behind the scenes (e.g. on-line advertising). Traditional auction theory has been focusing on single-dimensional mechanisms which usually allocate a product to a highest bidder or procure a service from the supplier with the lowest cost.

¹ This research is undertaken as part of the Danish Council on Strategic Research funded project on Foundations of Electronic Markets (CFEM). This is a collaborative project involving the Universities of Aarhus and Copenhagen, Copenhagen Business School, The Alexandra Institute, DCCA and industrial partners: DONG Energy, Energinet.dk, Inno:vasion, TradeExtension and Patrisia Market Design.

However, single-dimensional auctions focus solely on the price and hence fail to address issues related to the quality or the characteristics of the auctioned good. This is particularly important in cases where the auctioned good is not a well defined item, but instead it is a contract or a service in which cost or price is one of many parameters. For example, in a contract for the construction of a piece of public infrastructure, the price of the project is a very important parameter but it should not be the only one. In this case, the quality of the materials, the design and possible effects on the people and the environment are serious factors that should be taken into consideration when deciding the allocation¹ of the contract. Likewise, in a e-commerce application which for example may involve the purchase of web services, often parameters such as bandwidth, responsiveness, or even the reliability of the provider play a significant role with prospective clients not necessarily aiming for the cheapest service.

Against this background, Che in his seminal paper [1] has designed a series of auctions (first score, second score and second preferred score) to address the cases where the quality of a product is of equal importance to its cost. In these auctions, the suppliers report the quality of the item they are asked to provide and the costs involved in its production at the reported quality, which the mechanism maps into a single dimensional quantity, named as 'score'. All three auctions incentivise agents to truthfully report their costs, and result to an equal expected utility for the buyer, under the assumption that costs are independently distributed. This assumption was relaxed by Branco [2] who introduced a two-stage optimal multi-dimensional auction in a setting in which there was correlation among suppliers' costs. Bogetoft and Nielsen [5] propose more complex cost structures through the introduction of Data Envelopment Analysis (DEA [6, 7]) based competition. In this mechanism, and as opposed to second score auctions, all agents' bids determine the winner's payment.

Now, references to multi-dimensional auctions can be found in literature related to Computer Science and in particular multi-agent systems and e-commerce [8]. In particular, Bichler in [9] after comparing single and multi-dimensional auctions in an experimental web-based setting, showed that multi-dimensional auctions resulted in significant higher utility which paves the way for possible e-commerce applications. Furthermore, Beil and Wein [3] propose an iterative mechanism in which the buyer sequentially estimates each bidder's cost function through a series of score auctions. However, they assume that agents are truthful and hence do not address strategic behaviour. Parkes and Kalagnanam [4] propose an iterative multi-attribute procurement price-based auction in which suppliers at each round submit their bids and a winner which maximises the buyer's preference is selected. They show that their mechanism terminates with an outcome of modified Vickrey-Clarke-Groves allocation. Furthermore, multi-dimensional auctions can also be applied in settings where multiple suppliers can satisfy the principal's demand [10], although this exceeds the scope of this paper.

¹ A typical example is the European Union which has introduced a procurement directive based on two criteria: best economic value and lowest cost[11]. Examples include other government authorities, such as the US Department of Defense [1].

However, in all the above approaches there is a rather significant underlying assumption. It is assumed that the principal will enforce through the use of external and artificial means that the suppliers will report their truthful qualities. The above papers either do not include in their models the possibility of the participating agents having control on their reported qualities, or explicitly mention that the auction will be canceled, or an extremely heavy fine will be issued to the winner of the auction if the observed quality deviates from the reported one. We find this un-systematic approach to be unrealistic for most cases. For example, in state procurement auctions the government will face public scrutiny if an obviously incompetent supplier is selected and there may be a cost in the repetition of the auction, while in e-commerce a web content provider who fails to secure the required resource by its suppliers and consequently deliver it to its clients will lose credibility. In both cases, it is more important to create incentives that will motivate the suppliers to report their true quality and based on that for the buyer to decide, instead of canceling the auction or destroying the supplier in case there is a deviation between the reported and observed values.

It is this challenge we address in this paper through the design of two mechanisms, which initially allocate the contract to a specific supplier, while they guarantee that the suppliers reported their true qualities and costs through an additional payment after the contract is fulfilled and the outcome is observed. In more detail, in this paper we extend the state of the art in the following ways:

- We present a procurement mechanism in which the principal issues a two-step payment. The principal initially asks the agents to report their cost and quality, and based on these reports selects a single agent to allocate the contract of providing a service. The selected agent receives an initial payment based on its reports. After that agent fulfills its contract by providing a service at a specific quality which is observed by the principal, the principal issues an additional payment. The second payment balances between possible misreporting of quality in the beginning of the mechanism and the actual production in the final step of the mechanism.
- We introduce a second mechanism in which the principal issues an asymmetric second payment based only on the deviation between the reported and produced qualities. The agents producing more than they reported receive a positive payment, while those producing less than they reported receive a negative payment.
- We prove that for both our mechanisms truthful revelation of qualities and costs is a dominant and a weakly dominant strategy respectively and that they are immune to the combined misreporting of both qualities and costs. We also show that both mechanisms are individually rational. Finally, we show that both mechanisms result to optimal payments for the winning agent, equal to the optimal payment the agent receives in the standard second score auction. That is, the principal can elicit the agents' true qualities, without facing additional costs and without relying on artificial means (such as terminating the auction).

The rest of the paper is organised as follows: In Section 2 we introduce the notation and describe our setting in detail, while in Section 3 we discuss the tradi-

tional second score auction in order to provide some necessary background in multi-dimensional auctions. In Section 4 we introduce, analyse our mechanism and provide the proofs of its economic properties, while in Section 5 we introduce the asymmetric payments extension. Finally, in Section 6 we conclude and provide our future work.

1.2 The Context

We now introduce and describe our setting in more detail. We consider a principal interested in procuring a project, a service or an item from one of n rational and risk neutral agents, $i \in I = \{1, \dots, n\}$. The project depends on multiple parameters (e.g. bandwidth, latency, space in a data storage Cloud application). The parameters of a project undertaken by agent i are defined by an s -dimensional vector of qualities $y_i \in \mathbb{R}_0^s$. To simplify the analysis, and without loss of generality, we assume that for each agent the parameters of its service can be aggregated in one variable. Hence, each agent has a single quality profile denoted by y_i . In line with existing literature [1, 2], we assume that the agents have a prior knowledge of their quality, which as opposed to existing literature, may choose to misreport to the principal if they expect to increase their utility by doing so. The true quality is unknown to the principal, who does not have any means of verifying it before the fulfilment of the project. Therefore, the principal makes its decision based on the report of the agent.

Furthermore, agents are capable of producing different levels of qualities, but in order to produce the quality y_i agent i needs inputs. These inputs represent the cost in production and depend on each agent's efficiencies. The costs are private information to each agent and cannot be verified by any third party (i.e. the principal or other agents). The cost agent i faces in the production of its quality is modeled as a function of quality y_i and can be denoted as $x_i(y_i, c_i)$, where c_i represents the agent's private information about the cost of the project. We assume that the cost function is increasing in both quality and cost parameter and that is convex regarding the quality. These assumptions are realistic in all cases where we expect diminishing returns as the quality increases, for example due to limitations in production methods where significant resources are demanded for a small increase in production. Finally, while agents are aware of their cost parameters, the principal has only access to their distribution. We assume that c_i is independently and identically distributed over $[\underline{c}, \bar{c}]$ with $0 < \underline{c} < \bar{c} < +\infty$ according to a distribution with positive and continuously differentiable density function.

1.3 The Standard Second Score Auction

As previously mentioned, the second score auction has been widely used in order to procure goods or services while taking into consideration both price and quality, as opposed to standard auctions (i.e. English, Vickrey auctions) which focus solely on the prices. Despite this dual nature of the second score auctions, they only consider

agents' strategic behaviour with respect to their reported costs and not their reported qualities. They rely on the assumption that the auction will either be canceled, or the winner of the auction will face an extremely heavy fine if the promised quality is not delivered. Before we proceed to the analysis of our mechanisms which relax this assumption through a more rational and systematic approach, we provide the essential background for the understanding of the second score auction.

A second score auction maps the multi-dimensional bid (costs and qualities) into a single-dimensional variable (score). Effectively, for agent i : $(x_i, y_i) \mapsto S_i(x_i, y_i)$, with $S_i = V(y_i) - x_i(y_i, c_i)$, where $V(\cdot)$ is the principal's value function. The value function is modeled as an increasing and concave function of the quality. The interpretation of the above process is that an agent is willing to produce y_i if its paid at least x_i , with the principal's satisfaction is measured to be $V(y_i)$.

The mechanism proceeds as follows:

1. Principal asks n agents to submit their cost-quality bids (\hat{x}_i, \hat{y}_i) with $i \in \{1, \dots, n\}$
2. Each bid is assigned with a score $\hat{S}_i = S(\hat{x}_i, \hat{y}_i) = V(\hat{y}_i) - \hat{x}_i$.
3. The agent with the highest score wins the auction and is allocated the project².
4. The winner agent receives its payment by the principal: $\hat{P}_{(1)} = V(\hat{y}_{(1)}) - \hat{S}_{(2)}$, with subscripts (1) and (2) representing the winner and runner-up agent respectively.

In order to demonstrate how this mechanism works, let us consider a business owner who is interested in professional web presence. The candidate firms have received a list with the specifications she requires i.e. bandwidth, storage and CPU availability, access to multiple email accounts etc³. We assume that she has received three offers and that the corresponding bids are $(72.9, 9)$, $(16, 4)$, $(224, 16)$ where the first element is the provider's cost denoted by $x(y) = c_i * y^2$ with $c_i = \{0.9, 1, 0.875\}$. If her value function is $V(y) = 50 * \sqrt{y} + 25$ then the corresponding scores will be $S_i = \{102.1, 109, 1\}$. The second offer will be chosen and the payment will be equal to $125 - 102.1 = 22.9$.

The above mechanism has some very important properties. Most notably, truthful revelation of costs is a dominant strategy. Second, the mechanism is individually rational as it is easy to prove that the utility of the winner agent is always larger or equal (in case of tie-break) to 0. Finally, the optimal outcome maximises the socially efficient one, since the winner agent is the agent with the highest social welfare: $V(y_{(1)}) - x_{(1)}$ given that the mechanism is incentive compatible regarding cost and truthful reporting of the quality is assumed (the auction is canceled and a harsh penalty is issued to those who do not achieve the reported quality).

² Ties are ignored, although standard practices (i.e. random selection) may be applied. Che suggests in his paper that a random tie-breaking rule does not affect the expected payment in equilibrium [1].

³ We also assume that all these specifications can be aggregated into one quantity (referred as quality).

1.4 A Mechanism for Eliciting Both Costs and Outputs

We now relax the assumption regarding the guaranteed truthful reporting of qualities. In doing so, the principal faces the challenge of having to address an additional element of strategic behaviour on behalf of the agent. Indeed, in addition to misreporting its cost, now the agent can manipulate the principal by not delivering the promised quality.

We address the above challenge by designing a mechanism based on a two-step payment in order to create incentives which promote truthful reporting of both costs and qualities without relying on artificial means (i.e. cancellation of the auction). The proposed mechanism balances between two equally important concepts. That is, the agent must receive a reward for its actual production, while penalised for a possible misreporting of its costs and qualities.

Against this background, initially a standard second score auction is implemented in order to allocate the project to a specific agent who also receives its first payment. After the principal observes the actual outcome, secondary payment is issued to the selected agent. The first part of this additional payment consists of a symmetric penalty for the agent deviating from its reported quality, while in the second part it receives a compensation for its actual production.

In more detail the mechanism proceeds as follows:

1. Principal asks n agents to submit their cost-quality bids (\hat{x}_i, \hat{y}_i) with $i \in \{1, \dots, n\}$.
2. Each bid is assigned with a score $\hat{S}_i = S(\hat{x}_i, \hat{y}_i) = V(\hat{y}_i) - \hat{x}_i$.
3. The agent with the highest score wins the auction and is allocated the project.
4. The winner agent receives its **first payment** by the principal: $\hat{P}_{(1)} = V(\hat{y}_{(1)}) - \hat{S}_{(2)}$.
5. Winning agent produces quality $y_{(1)}$ (not necessarily equal to $\hat{y}_{(1)}$).
6. Principal observes winning agent's quality production and issues the **second payment**:

$$B(\hat{y}_{(1)}, y_1) = d(\hat{y}_{(1)}, y_1) \times [V(\hat{y}_{(1)}) - \hat{S}_{(2)}] + [V(y_{(1)}) - \hat{S}_{(2)}] \quad (1.1)$$

where $d(\hat{y}_{(1)}, y_1)$ is the function that measures the effect of winner's misreporting and consequently the amount of the initial payment the agent has to return back. In the following section we discuss the properties such a function should have and we proceed to define it.

1.4.1 Definition of $d()$ function

The secondary payment consists of two parts. In the first part the principal penalises the selected agent for a possible deviation between the reported and produced quality, while in the second part the selected agent receives a payment for its actual production. The $d()$ function is designed to guarantee that the selected agent will be penalised for not being able to meet its initial claims. The penalty is symmetric, that

is the selected agent is penalised for its misreporting even if its misreporting results to an increase of the principal's value (through over-production). This will be necessary for applications in which over-production has negative effect. For example, over-production may be unwanted due to limitations in storage or transport capacity, or market conditions (over-supply). Alternatively, there may be more complex cases where agents may have to provide costly estimates of their qualities. In these cases, it will be unacceptable if the same amount of misreporting resulted to a different type of payment (i.e. penalty vs reward).

Therefore, the properties we seek in $d()$ function are the following:

1. The first part of the payment must be a penalty, hence $d()$ must be negative for every $\hat{y}_{(1)} \in [0, +\infty)$.
2. The penalty must be proportional of the misreporting, that is minor deviation should result to a minor penalty, with the penalty increasing as misreporting increases.
3. There should be no distinction between agents which under or over reported their qualities, hence function $d()$ must be symmetrical.
4. The secondary payment must be maximised at truthful reporting $\hat{y}_{(1)} = y_{(1)}$ in order to guarantee incentives.

A function which satisfies the above assumption is the following⁴:

$$d(\hat{y}_{(1)}, y_{(1)}) = - \left[\frac{(\hat{y}_{(1)} - y_{(1)})^2}{y_{(1)}^2} \right] - 1 \quad (1.2)$$

1.4.2 Proof of Economic Properties

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

1. Truthful revelation of agents' qualities is a dominant strategy.
2. Truthful revelation of agents' costs is a weakly dominant strategy.
3. The mechanism is immune to combined strategic behaviour of both qualities and costs.
4. The mechanism is individually rational.
5. The optimal payment is equal to the optimal payment of a standard second score auction mechanism.

Lemma 1 *Truthful revelation of the agents' qualities is a dominant strategy given that they truthfully reported their costs.*

⁴ We are aware that the choice of the $d()$ function appears to be arbitrary. It should be noted that a more complicated analysis of the function exceeds the scope of this paper. This analysis will be central in future work where we consider agents facing costs in the process of identifying their qualities.

Proof. We prove that truthful revelation of agents' qualities is a dominant strategy by showing that winner's utility is maximised at truthful reporting (i.e. $\hat{y}_{(1)} = y_{(1)}$).

The winning⁵ agent's utility⁶ is equal to:

$$U(\hat{y}) = V(\hat{y}) - S_2 + d(\hat{y})[V(\hat{y}) - S_2] + V(y) - S_2 - x(y) \quad (1.3)$$

With its derivative being equal to:

$$U'(\hat{y}) = [1 + d(\hat{y})]V'(\hat{y}) + d'(\hat{y})[V(\hat{y}) - S_2] \quad (1.4)$$

For $\hat{y} = y$ and given that $d(y) = -1$ and $d'(y) = 0$ the derivative is equal to 0.

The second derivative for the utility function is:

$$U''(\hat{y}) = d'(\hat{y})V'(\hat{y}) + [1 + d(\hat{y})]V''(\hat{y}) + d''(\hat{y})[V(\hat{y}) - S_2] + d'(\hat{y})V'(\hat{y}) \quad (1.5)$$

For $\hat{y} = y$ and given that $d(y) = -1$ and $d'(y) = 0$ the second derivative is equal to $U''(y) = d''(y)[V(y) - S_2]$ which is negative given that $d''(y) < 0$ and $V(y) - S_2 > 0$. Regarding $V(y) - S_2$, for a selected agent who has truthfully reported its cost: $S_1 = V(y) - x(y) > S_2 \Rightarrow V(y) > S_2$. Finally, the selected agent's utility is equal to $S_1 - S_2$ which positive (unless there is a tie between the first and second agent's scores). \square

Lemma 2 *Truthful revelation of the agents' costs is a weakly dominant strategy given that they truthfully reported their qualities.*

Proof. We prove this by contradiction. Let x and y be an agent's true cost and quality, and S the score that corresponds to these true values, while $\hat{x}, \hat{y}, \hat{S}$ the reported ones. Furthermore, let x_2, y_2, S_2 be the bids and score of the runner up agent (i.e. $\hat{S} > S_2$).

First, let the agent's misreporting have an effect 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 its cost s.t. $\hat{S} > S_2$ given that $S < S_2$. The agent achieves this by reporting a lower cost than its actual one i.e. $\hat{x} < x$. Under optimal reporting of quality, the utility of an agent misreporting its cost in Case (1) will be negative i.e. $U(y) = V(y) - x(y) - S_2 = S - S_2 < 0$.
 - In Case (2) agent reports its cost s.t. $\hat{S} < S_2$ given that $S > S_2$. The agent would have won the auction, but instead reports a cost greater than its actual one i.e. $\hat{x} > x$. As a result, the agent loses the auction and consequently receives zero utility.

⁵ From this point we omit subscript (1) for notational convenience

⁶ After the actual quality y has been observed, function $d(\hat{y}, y)$ can be simplified to $d(\hat{y})$, since \hat{y} is the only unknown variable.

Second, we assume that the agent misreports its cost 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 zero. Had the agent already won the auction, misreporting would not result to additional benefits. Specifically, both payments depend on the second lower score and the reported and actually produced (for the second stage) quality. \square

Theorem 1 *The mechanism is immune to combined misreporting of both qualities and costs.*

Proof. In the above proofs we showed that truthful reporting of quality is an optimal strategy if the agent reports truthfully its cost, and that the same holds for an agent's costs, given that it truthfully reported its quality. However, given the multi-dimensional nature of the bids an agent could attempt to manipulate the principal by misreporting both costs and qualities. Specifically, we consider these two general cases:

1. Agent wins the auction by misreporting both its quality and cost
 2. Agent wins the auction with the misreporting having no effect on the auction's outcome
 3. Agent loses the auction despite or due to its misreporting
- In Case (1) the agent reports its quality and cost s.t. $\hat{S} > S_2$, while $S < S_2$. For example, an agent may be inclined to report both lower cost and higher quality than its true values in order to guarantee that it will win the auction. Let an agent's reported quality be $\hat{y} = ky$ with $k \in [0, +\infty)$ being the coefficient of under-reporting ($k \leq 1$) or over-reporting ($k > 1$), while the reported cost remains \hat{x} . The utility derived by that agent is:

$$U = V(ky) - S_2 - [(k - 1)^2 + 1][V(ky) - S_2] + V(y) - S_2 - x(y) \Rightarrow$$

$$U = -(k - 1)^2[V(ky) - S_2] + S_1 - S_2$$

Now since the agent misreported both its cost and quality and won the auction while it should not have won $k \neq 1, S < S_2$ and $\hat{S} > S_2$ the above expression is negative given that $\hat{S} = V(ky) - \hat{x}(ky) > S_2 \Rightarrow V(ky) > S_2$. This type of combined misreporting leads to negative utility.

- 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. In this case, $S > S_2$, and although it is possible that $U > 0$, we know from Lemma 1 that the agent's utility is maximised at $\hat{y} = y$ if $V(y) - S_2 > 0$. Now, given that in this case the agent has already won the auction: $S - S_2 = V(y) - x(y) - S_2 > 0 \Rightarrow V(y) - S_2 > 0$. So misreporting would deviate from the optimal outcome.

- In Case (3) the agent does not win the auction after misreporting its cost and quality. This may have happened as a result of the misreporting (the agent would have won otherwise), or despite the misreporting. In both those cases, the utility the agent will receive will be equal to zero.

We have showed that combined misreporting of both costs and qualities leads either to negative utility or to a non-optimal outcome, hence the mechanism is immune to this type of strategic behaviour. \square

Theorem 2 *The mechanism is individual rational*

Proof. We have shown in the above proofs that for truthful reporting of qualities the selected agent's utility is equal to $V(y) - S_2 - x(y)$, which is positive for truthful reporting of costs. This is also the optimal payment the agent receives for reporting its true cost and quality in the standard second score auction. It can be seen, that the principal can elicit agents' true quality without additional cost on its behalf and without relying on artificial means such as cancellation of the auction. \square

Corollary 1 *The agent's optimal payment for reporting its true cost and quality is equal to the payment issued in a standard second score auction.*

1.5 A Mechanism with asymmetric secondary payment

In the previous section, we introduced a mechanism in which the principal penalises the winning agent for failing to provide the quality it reported. The principal does not discriminate among those agents who under or over reported and balances the fine with a fixed payment equal to the payment the agent would have received in the original second score auction. We focused on the case where the principal's primary focus is the quality's misreporting even if it ends up receiving value higher than it initially expected. Although this may be necessary in several cases (i.e. limited storage capacity, flooding of market through over-supply) or in more complex settings (i.e. costly estimates of qualities), it is natural to assume that in a setting with no special conditions or uncertainty regarding the quality, over-production should not be penalised.

Against this background, in this section we introduce a mechanism in which the principal does not directly penalise those agents who end up producing more than they reported without necessarily incentivising them to do so. The principal still values truthful reporting the most, however it is not willing to equally penalise all agents to achieve it. The first payment for this mechanism is identical to the first one, while the secondary payment issued after the actual production, rewards a selected agent which has produced a good or a service of quality higher than the one it reported initially in the mechanism, while it issues a fine to a selected agent with the exact opposite behaviour.

In more detail the payment issued by the principal after observing the agent's production $y_{(1)}$ (not necessarily equal to $\hat{y}_{(1)}$), is the following:

$$B(\hat{y}_{(1)}, y_1) = d(\hat{y}_{(1)}, y_1) \times [V(y_1) - V(\hat{y}_{(1)})] \quad (1.6)$$

where $d(\hat{y}_{(1)}, y_1)$ is the function that measures the effect of winner's misreporting. As opposed to Mechanism 1, in this mechanism function $d()$ may reward an agent even if it is misreporting its quality. In the following section we discuss the design of this function and how the additional payment can be balanced in order to guarantee incentives in this mechanism.

1.5.1 Definition of $d()$ function

In this mechanism, the winner of the auction is not necessarily directly penalised for choosing to produce a good at a different quality than the promised one. In this case the selected agent will not be asked to return a part of its initial payment if the principal has benefited from the agent's misreporting. We take this into consideration by designing $d()$ function so that an agent will not be directly penalised for causing an increase in principal's value (through over producing), although it may still be far from doing the best it can. Hence, it is important to differentiate between the misreporting that inflicts losses and that one which causes gains for the principal. In order to do so, function $d()$ must have the following properties:

- **For the agent who reported less than it produced and thus created a surplus,**
 $\hat{y} \in [0, y]$:
 1. Given that $V(y) - V(\hat{y}) > 0$, $d()$ must be positive for every $\hat{y} \in [0, y]$ so the secondary payment is in fact a reward.
 2. The agent should receive a portion of the principal's surplus $V(y) - V(\hat{y})$.
 3. The agent should not be inclined to under-report systematically, hence it should receive a small portion of that surplus for significant under-reporting.

** Hence, for $\hat{y} \in [0, y]$, $d(\hat{y})$ should be an increasing and positive function taking values in $[0, 1]$.
- **For the agent who reported more than it produced and thus created a deficit,**
 $\hat{y} \in (y, +\infty)$:
 1. Given that $V(y) - V(\hat{y}) < 0$, $d()$ must be positive for every $\hat{y} \in (y, +\infty)$ so the secondary payment is in fact a penalty.
 2. The agent's penalty should increase as the principal's deficit in quality $V(y) - V(\hat{y})$ increases. The agent possibly manipulated the first payment of the mechanism by reporting high quality, but it turned out it could not deliver it.
 3. The agent should not be over-penalised for a small amount of misreporting, hence it should receive a fine proportional to the deficit of principal's value.

** Hence, for $\hat{y} \in (y, +\infty)$, $d(\hat{y})$ should be an increasing and positive function taking values in $(1, +\infty)$.

A function which satisfies all the above properties is:

$$d(\hat{y}_{(1)}, y_{(1)}) = \frac{(\hat{y}_{(1)} - y_{(1)})^3}{y_{(1)}^3} + 1 \quad (1.7)$$

1.5.2 Proof of Economic Properties

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

1. Truthful revelation of agents' qualities is a dominant strategy.
2. Truthful revelation of agents' costs is a weakly dominant strategy.
3. The mechanism is immune to combined strategic behaviour of both qualities and costs.
4. The mechanism is individually rational.
5. The optimal payment is equal to the optimal payment of a standard second score auction mechanism.

Lemma 3 *Truthful revelation of agents' qualities is a dominant strategy, given that they truthfully reported their costs.*

Proof. We prove that truthful revelation of agents' qualities is a dominant strategy by showing that winner's utility is maximised at truthful reporting (i.e. $\hat{y}_{(1)} = y_{(1)}$) or after simplifying the notation $\hat{y} = y$.

The selected agent's utility after the a project of quality y has been produced and observed, is given by:

$$U(\hat{y}) = V(\hat{y}) - S_2 + d(\hat{y})[V(y) - V(\hat{y})] - x(y) \quad (1.8)$$

With its derivative being equal to:

$$U'(\hat{y}) = [1 - d(\hat{y})]V'(\hat{y}) + d'(\hat{y})[V(y) - V(\hat{y})] \quad (1.9)$$

It is clear, that for $\hat{y} = y$ and given that $d(y) = 1$ and $d'(y) = 0$ the derivative is equal to 0.

The second derivative of the utility function is:

$$U''(\hat{y}) = [1 - d(\hat{y})]V''(\hat{y}) + d''(\hat{y})[V(y) - V(\hat{y})] - 2d'(\hat{y})V'(\hat{y}) \quad (1.10)$$

For $\hat{y} = y$ and given that $d(y) = 1$ and $d'(y) = d''(y) = 0$ the second derivative is equal to 0], hence the second derivative test is inconclusive for the stationary point y . We will show that y is a maximum point by examining the sign of the derivative of the utility function for $y \in [0, +\infty)$. Indeed, for Equation 1.9 it can be said that it is positive for $\hat{y} \in [0, y]$, while it is negative for $\hat{y} \in (y, +\infty)$. Hence the utility function is increasing in $[0, y]$, while it is decreasing in $(y, +\infty)$. That is $\hat{y} = y$ is a maximum point with $U(y) = S_1 - S_2 > 0$ (unless there is a tie), hence truthful revelation of agents' reported quality is a dominant strategy. \square

Lemma 4 *Truthful revelation of agents' reported costs is a weakly dominant strategy given that they truthfully reports their qualities.*

Proof. The proof for this theorem is identical to the proof of Lemma 2. Despite the fact that in this mechanism there is a different secondary payment, under optimal reporting of quality, the agent derives an identical utility, equal to $V(y) - S_2 - x(y) = S - S_2$, where S_t is the score that corresponds to truthful reporting of quality and its cost. \square

Theorem 3 *The mechanism is immune to combined misreporting of both qualities and costs.*

Proof. In the above proofs we showed that truthful reporting of quality is an optimal strategy if the agent reports truthfully its cost, and that the same holds for an agent's costs, given that it truthfully reported its quality. However, given the multi-dimensional nature of the bids an agent could attempt to manipulate the principal by misreporting both costs and qualities. An additional complication appears given that there may be a positive secondary payment for a $\hat{y}^* \in [0, y]$, and hence it is important to determine whether it is possible for an agent to manipulate the reporting of costs early in the mechanism in order to guarantee positive secondary payment. We still consider the two following general cases:

- Agent wins the auction by misreporting both its quality and its cost
- Agent wins the auction with the misreporting having no effect on the auction's outcome.
- Agent loses the auction despite or due to its misreporting
- In Case (1) agent reports both its quality and cost s.t. $\hat{S} > S_2$, while $S < S_2$. Let $y^* = k^*y$, with k^* being the coefficient of under-reporting ($k^* \leq 1$) or over-reporting ($k^* > 1$), with k^*y maximising the secondary payment $B(\hat{y}, y)$ (but not the overall agent's utility). The utility derived by that agent is:

$$U = V(k^*y) - S_2 + [(k^* - 1)^3 + 1][V(y) - V(k^*y)] - x(y) \Rightarrow$$

$$U = [S - S_2] + [(k^* - 1)^3[V(y) - V(k^*y)]]$$

Now since the agent misreported both its cost and quality and won the auction while it should not have won the above expression is negative. Indeed $S - S_2 < 0$ and $[(k^* - 1)^3[V(y) - V(k^*y)]] < 0$ for both $k^* < 1$ and $k^* > 1$

- In Case (2) the agent would have won the auction anyway and its possible misreporting of quality and cost had no impact on that outcome. It is important now to examine whether it is still in an agents best interest to report y^* . In this case, $S - S_2 > 0$, and $[(k^* - 1)^3[V(y) - V(k^*y)]] < 0$. Although, it may be possible that $U > 0$, we know from Lemma 3 that the utility will be maximised at $k = 1$. Hence, it will not be optimal for an agent not to report its true quality, given that it will end up receiving less for misreporting.

- In Case (3) the agent does not win the auction after misreporting its cost and quality. This loss may have happened due or despite the agent's misreporting and a result of it the utility it will receive will be equal to zero.

We have showed that combined misreporting of both costs and qualities leads either to negative utility or to a non-optimal outcome, hence the mechanism is immune to this type of strategic behaviour. \square

Theorem 4 *The mechanism is individual rational*

Proof. We have shown in the above proofs that for truthful reporting of qualities the selected agent's utility is equal to $V(y) - S_2 - x(y)$, which is positive for truthful reporting of costs. This is also the optimal payment the agent receives for reporting its true cost and quality in the standard second score auction. It can be seen, that the principal can elicit agents' true quality without additional cost on its behalf and without relying on artificial means such as cancellation of the auction. \square

Corollary 2 *The agent's optimal payment for reporting its true cost and quality is equal to the payment issued in a standard second score auction.*

1.6 Conclusion

In this paper we introduced two novel multi-dimensional mechanisms based on a two-step payment in a setting in which the principal cannot enforce truthful reporting of agents' qualities through artificial means such as the cancellation of the auction or extremely high fines. Initially, the principal procures an item or a service from the winner of a standard second score auction and it issues a second score auction payment based on that agent's reported quality and cost. That agent then fulfils its part of the contract and produces an output of a certain quality. The principal, after having observed the production, calculates a second payment based on both the selected agent's reported and produced quality.

Regarding the secondary payment, in the first mechanism it consists of two parts. In the first part the principal issues a symmetric penalty to the selected agent for misreporting its quality, while in the second part it compensates that agent through a standard second score auction payment based on the agent's produced quality and not its reported. In the second mechanism, the principal issues an asymmetric payment to the selected agent. That is, it rewards that agent if the procured service is of higher quality to the reported one, while it penalises it for the exact opposite. Despite the different approaches, under truthful reporting both mechanisms result to the same total payment which is equal to payment issued to the winner of the standard second auction. Hence, there is no cost in the principal's attempt to guarantee incentives regarding the agents' reported qualities, and therefore no reason at all for enforcing them upon the agents.

For both mechanisms we provided proofs of their economic properties. Specifically, we showed that they are incentive compatible regarding the reported qualities

(dominant strategy) and costs (weakly dominant strategy) and individually rational. In addition to this, we showed that they are immune to agents' combined misreporting of their parameters.

In our future work we intend to consider a setting in which the agents are not aware of their specific quality. Agents will only have a probabilistic estimate of their production and based on that estimate the principal will have to make a decision. Agents will face costs in the generation of their estimates and hence they will have to balance between the amount of resources they will invest on generating their estimate and the consequences they may face for providing imprecise or inaccurate estimates. By addressing these issues, we will be able to apply this mechanism in more complex and dynamical domains where there is uncertainty regarding the production. A typical example is the production of electricity through specific renewable sources which depend on weather conditions (i.e. solar, wind power) and hence subject to forecasting. In this case the principal will have to elicit truthful and precise estimates, while the focus of the mechanism remains on the actual production. Also, costs involved in both the estimate and production will have to be balanced so that the whole process is a viable option for both the principal and the participating agents.

References

1. Che, Y-K.: Design competition through multidimensional auctions. *RAND Journal of Economics* **244** (1993) pp. 668–680.
2. Branco, F.: The design of multidimensional auctions. *RAND Journal of Economics* **281** (1997) pp. 63–81.
3. Beil, D. R. and Wein, L.: An inverse-optimization-based auction mechanism to support a multiattribute RFQ process. *Management Science* **49**11 (2003) pp. 1529–1545.
4. Parkes, D. and Kalagnanam, J.: Models for Iterative Multiattribute Procurement Auctions. *Management Science* **51**3 (2005) pp. 435–451.
5. Bogetoft, P. and Nielsen, K.: DEA based auctions. *European Journal of Operational Research* **184**2 (2008) pp. 685–700.
6. Charnes, A., Cooper, W. W. and Rhodes, E.: Measuring the efficiency of decision making units. *European Journal of Operational Research* **26** (1978) pp. 429–444.
7. Charnes, A., Cooper, W. W. and Rhodes, E.: Short Communication: Measuring the efficiency of decision making units. *European Journal of Operational Research* **3** (1979) pp. 339.
8. He, M., Jennings, N. R. and Leung, H.-F.: On Agent-Mediated Electronic Commerce *IEEE Transactions on knowledge and data engineering* **15**4 (2003) pp. 985–1002.
9. Bichler, M.: An experimental analysis of multi-attribute auctions. *Decision Support Systems* **29** (2000) pp. 249–268.
10. Bichler, M. and Kalagnanam, J.: Configurable Offers and Winner Determination in Multi-Attribute Auctions *European Journal of Operational Research* **160** (2003) pp. 380–394.
11. Asker, J. and Cantillon, E.: Properties of scoring auctions. *RAND Journal of Economics* **39**1 (2008) pp. 69–85.