Auctioning Incentive Contracts:
An Experimental Study

Sander Onderstal[1]
Arthur Van de Meerendonk[2]

[1] University of Amsterdam, and Tinbergen Institute;
[2] University of Maastricht.

Auctioning Incentive Contracts: An Experimental Study

Sander Onderstal

Amsterdam School of Economics, University of Amsterdam, the Netherlands, onderstal@uva.nl

Arthur Van de Meerendonk

Maastricht School of Governance, University of Maastricht, the Netherlands, meerendonk@M13.nl

In this note, we experimentally examine the relative performance of price-only auctions and multi-attribute auctions. We do so in procurement settings where the buyer can give the winning bidder incentives to exert effort on non-price dimensions after the auction. Both auctions theoretically implement the surplus maximizing mechanism. Our experiment confirms this result. Moreover, we observe that the "pie" is shared the same in both auctions between buyer and suppliers both in theory and in the lab (after accounting for learning effects).

Key words: Procurement; Price-only auctions; Multi-attribute auctions; Incentive Contracts; Laboratory Experiment

1. Introduction

In procurement contests, buyers typically have to decide whether to let potential suppliers compete on just one dimension (usually the price) or on several (such as price, quantity, and quality). Both ways of procurement are common in practice. In a price-only (PO) auction, the buyer only invites bids on the price and usually the winner is the supplier offering the lowest price. In contrast, in a multi-attribute (MA) auction, suppliers submit bids on several dimensions. Practical applications of MA auctions include tenders for weapons systems (Che 1993), B2B internet auctions (Burmeister et al. 2002), highway construction procurement (Bajari and Lewis 2009), and welfare-to-work tenders (Onderstal 2009).

The co-existence of these mechanisms gives rise to the following questions:

- Which mechanism generates the highest utility for the buyer?

- Which mechanism generates the highest profit for the suppliers?

- Which mechanism is the most efficient, i.e., generates the highest total surplus?

In this paper, we aim to answer the above questions in a laboratory experiment. These questions are not only of academic interest. Procurement is big business, both in the public (B2G) and in the private do-

main (B2B). In 2002, total expenses on public procurement in the European Union amounted to roughly €1500 billion, i.e., 16% of GDP (European Commission 2004). In the same year, DaimlerChrysler handled a procurement volume of more than 10 billion Euros (Burmeister et al. 2002). Additionally, the buyer may wish to incentivize the winning supplier to exert costly effort after winning the procurement contest. In an MA auction, the buyer can do so by letting the suppliers also submit bids on relevant non-price attributes. In a PO auction, the buyer can award the winning supplier an incentive contract which aims to make the payments to the supplier depend on the supplier's effort on non-price dimensions. Indeed, incentive contracts are widely used in procurement practice. For instance, Italian law requires public procurement contracts to include a penalty fee between 0.03% and 0.1% of the total contract price for each day of delay in completion of the contracted works (D'Alpaos et al. 2009). Similar figures apply for highway construction in the US (Herbsman et al. 1995). The questions that are raised above are also interesting in the light of the issue whether public bodies should procure on the basis of "the lowest price offered", i.e., a PO auction, or on the basis of "the economically most advantageous tender" (EMAT), i.e., an MA auction. According to the European Commission (2003), "[p]rocurement systems should be used that ensure that enterprises bear some responsibility for quality and performance of the completed works, which would lead to the EMAT." In this note, we will show that PO auctions with a well designed incentive contract may perform just as well as MA auctions.

The structure of our setting is as follows. Suppliers compete to win a project in either a PO auction or an MA auction. In the PO auction, the highest bidder wins the project and is rewarded a fee for each unit of output she produces in the project. In the MA auction, each supplier submits bids on two dimensions: the "output" she will deliver if she wins the project and the payment she wishes to receive in return. Output is broadly interpreted: cost reduction, delivery time, quality, quantity, and so forth. The buyer assigns the project to the supplier whose bid results in the highest non-negative score (a well-defined function of these bids). The winning supplier obtains her monetary bid plus a bonus [minus a malus] if her actual output is more [less] than the output in her bid. We will answer the above questions in two environments: a deterministic one, in which output equals effort, and a stochastic one, in which output equals effort plus

a disturbance term. The disturbance term could capture the business cycle, uncertainty about the difficulty of completing the project, and so forth.

Our main results are the following. In theory, the most efficient supplier always wins in equilibrium in both the PO auction and the MA auction. Moreover, the auctions are outcome equivalent in the sense that both generate the same utility for the buyer, the same profit for the suppliers, and the same total surplus. Finally, both mechanisms can implement the surplus maximizing mechanism. In our experiment, we observe the following. Both the PO auction and the MA auction select the most efficient supplier in about 3 out of 4 times. Still, total surplus is very close to the theoretical optimum. Initially, the MA auction yields higher average utility for the buyer than the PO auction, while the opposite holds true for the average supplier profit. This, if true, would support a preference for procuring organizations towards MA auctions over PO auctions. However, these differences vanish after subjects have had the opportunity to learn. Our findings in the deterministic and stochastic environments are similar.

From the theoretical literature, the following results are relevant for our study. McAfee and McMillan (1986, 1987) and Laffont and Tirole (1987) construct optimal auctions of incentive contracts for settings that include ours. An MA auction can implement the optimal mechanism in the case of independent efficiency levels (Che 1993, Branco 1997) in contrast to PO auctions where the non-price attributes are fixed ex ante (Asker and Cantillon 2008). The reason that PO auctions with pre-determined non-price attributes are not optimal is that in the optimal contract, the optimal non-price attributes depends on the winner's efficiency level.

We are not the first to study the relative performance of PO and MA auctions in a laboratory experiment. Experiments by Bichler (2000) and Chen-Ritzo et al. (2005) indicate that MA auctions outperform PO auctions in terms of efficiency, buyer utility, and supplier profits. However, subjects were not incentivized to provide more than the minimum effort on non-price dimensions. As said, if this is the case, (optimal) MA auctions outperform PO auctions in theory (Asker and Cantillon 2008). Engelbrecht-Wiggans et al.'s (2007) experiment provides another means of comparing the two mechanisms. In their setting, a supplier's quality level is fixed and correlated with her cost level. In theory, the PO auction dominates the

MA auction in terms of buyer surplus if and only if the number of bidders is sufficiently large. The experiment confirms this theoretical result.[1] The main difference between these experiments and ours is that we incentivize suppliers to exert effort on non-price attributes ("output" in our model). In the MA auction, output is one of the bidding dimensions, while in the PO auction, we give the winning supplier ex post incentives to expend output increasing effort. We have chosen our parameters such that in theory, both mechanisms perform equally well.

In contrast to MA auctions, PO auctions have received ample attention in the experimental literature in settings where only the price is relevant (Kagel 1995 provides an overview). One important insight from this literature is that mechanisms that are in theory outcome equivalent may yield different outcomes in the lab. As said, in our setting, the PO auction and the MA auction are outcome equivalent. So, the experiments on PO auctions indicate that a priori it may still be the case that one mechanism outperforms the other.

The set-up of the remainder of this note is as follows. Section 2 includes the theoretical predictions. In Section 3, we describe our experimental design. In Sections 4 and 5, we discuss the experimental results on an aggregate level and on an individual level respectively. Section 6 contains a conclusion. Proofs of propositions are relegated to the appendix.

2. Theory

A risk neutral buyer procures an indivisible project. We assume that n risk neutral suppliers participate in the procurement. Each supplier i, $i = 1, ..., n$, upon winning the project, is able to expend effort e_i at cost $C_i = C(e_i, t_i)$ where t_i is supplier i's efficiency level.[2] We assume that $C(e_i, t_i) = e_i^2/(2ct_i)$ where $c > 0$ is a

---

[1] Haruvy and Katok (2007) further explore MA auctions in the setting of Engelbrecht-Wiggans et al. (2007). In their experiment, they find that sealed-bid MA auctions outperform dynamic ones, especially in situations where suppliers have information about each others' quality.

[2] Note that we assume one-dimensional types. Asker and Cantillon (2008) show that this is not necessarily with loss of generality in that a vector of types can be substituted by a one-dimensional "pseudo-type".

scaling parameter. The suppliers draw the t_i's independently from the uniform distribution on the interval $[a, b]$, $b > a > 0$.

If supplier i wins, she obtains utility $U_i = T_i - C_i$ where T_i is the monetary transfer that she receives from the buyer. We normalize her utility to zero if she does not win the project. The output m_i of supplier i if she wins the project equals $m_i = e_i + \varepsilon$, where ε represents "the state of the economy". We assume that ε is a random variable with mean zero and finite variance and that the winner is only informed about ε after she has chosen her effort level. The buyer only observes m_i, i.e., he cannot observe e_i or ε. The buyer's utility equals $U_B = sm_i - T_i$, where $s > 0$ denotes the marginal benefits for the buyer of the supplier's output. Note that $\mathrm{E}\{m_i\} = e_i$, so that a risk neutral buyer aims at maximizing $\mathrm{E}\{U_B\} = se_i - T_i$.

Let S denote total surplus of the project. Total surplus is given by $S = U_B + \sum_i U_i$, which implies that transfers between the buyer and a supplier are welfare neutral. A mechanism is called efficient if S is maximized. A socially optimal mechanism maximizes expected surplus under the restriction that the suppliers play a Bayesian Nash equilibrium, and under a participation constraint (each participating supplier must at least receive zero expected utility). Because transfers from the buyer to the suppliers are welfare neutral, the first-best socially optimal mechanism can be readily constructed. First of all, the buyer selects the most efficient supplier, i.e., the supplier with the highest type t_i. Secondly, the buyer induces this supplier to exert effort $e^*(t_i)$ for which

$$e^*(t_i) = cst_i. \tag{1}$$

We assume that $\varepsilon > - csa$, so that for the optimal output it holds true that

$$m_i = e^*(t_i) + \varepsilon \geq e^*(a) + \varepsilon = csa + \varepsilon > 0.$$

Finally, the buyer pays a transfer to the suppliers such that all suppliers get at least zero expected utility. The buyer can implement the socially optimal mechanism using a PO auction. In the first period, suppliers independently submit bids in the first-price sealed-bid auction. The highest bidder wins the project, and pays her bid. In the second period, the highest bidder exerts effort in the project and receives $T = m\tau_p - B$ if she bids B and her output equals m, where $\tau_p > 0$ is the per unit transfer to the winning supplier. The following proposition establishes that the PO auction is efficient if $\tau_p = s$.

PROPOSITION 1. *In the unique subgame perfect Bayesian Nash-equilibrium of the PO auction, a supplier with efficiency parameter t bids*

$$B^{PO}(t) = (1/2)(\tau_p)^2[t - (t - a)/n].$$

If she wins the project, her effort level will be

$$e^{PO}(t) = ct\tau_p.$$

The most efficient supplier wins the project. The PO auction implements the socially optimal mechanism if $\tau_p = s$.

Che (1993) shows that also an MA auction can implement the socially optimal mechanism. In an MA auction, bidders submit bids on several dimensions. The auctioneer assigns the project to the bidder submitting the highest non-negative score (a well-defined function of these bids). In our setting, these dimensions are the supplier's (promised) effort e and payment T. The scoring function is $\sigma(e, T) = \tau_m e - T$, where $\tau_m > 0$. The project will be assigned to the supplier with the highest score, unless the score is negative. In the latter case, the project will not be assigned. If the output of the winning supplier is above [below] the effort in her bid, she will get a bonus [malus] equal to τ_m for each unit of output above [below] her effort level. Proposition 2 establishes that the MA auction implements the socially optimal mechanism if $\tau_m = s$.

PROPOSITION 2. *In the Bayesian equilibrium of the MA auction, a supplier with efficiency parameter t bids*

$$\{e^{MA}(t), T^{MA}(t)\} = \{ct\tau_m, (1/2)(\tau_m)^2[t + (t - a)/n]\}.$$

The most efficient supplier wins the project. The MA auction implements the socially optimal mechanism if $\tau_m = s$.

The following result follows immediately from Propositions 1 and 2. The observation that the PO auction and the MA auction are utility equivalent if $\tau_m = s$ follows from the facts that both mechanisms are

efficient, $e^{PO}(t) = e^{MA}(t)$, and $se^{PO}(t) - B^{PO}(t) = T^{MA}(t)$, so that effort and the net payments in both mechanisms coincide.

PROPOSITION 3. *If $\tau_p = \tau_m = s$, the PO auction and the MA auction are efficient and generate the same expected utility for the buyer and the suppliers.*

3. Experimental Design

In a computerized laboratory experiment, we studied the PO auction and the MA auction in a setting that is closely related to our theoretical model.[1] We used the following parameters: $n = 2$, $[a, b] = [50,100]$, $c = 1/50$, and $\tau_p = \tau_m = s = 10$. The two main differences between the model and the experimental design are the following. First, the efficiency parameters were drawn according to a uniform distribution on a grid between 50 and 100 with 1 as the smallest step instead of from the entire interval. Second, we approximated the continuous bidding and effort spaces with grids with 1 as the smallest step. The experiment was conducted at the University of Amsterdam in May 2007.

The experiment consisted of two treatments, one for each auction. We held four sessions, two sessions for each treatment, resulting in a between-subject design. In all sessions, the subjects were separated in groups of four. A session consisted of four parts. In the first two parts, all subjects obtained a project and were paid ten points for each unit of output. These parts can be seen as practice rounds in which subjects could learn to find the optimal effort in a project and to calculate the profitability of a project. In parts 3 and 4, subjects competed in one of the two auctions. Parts 1 and 2 [3 and 4] consisted of 5 [10] rounds each. In parts 1 and 3, we confronted subjects with a deterministic environment in which $P\{\varepsilon = 0\} = 1$, while in parts 2 and 4, subjects competed in a stochastic environment in which $P\{\varepsilon = -2\} = P\{\varepsilon = 0\} = P\{\varepsilon = 2\} = 1/3$. At the start of the session, the 24 [36] participants in the PO [MA] auction were clustered in groups of four. These groups were maintained throughout the session. Before the start of each round,

---

[1] A translation of the instructions can be found in the electronic companion to this paper.

the subjects were randomly re-matched to another player in their group of four, resulting in 6 [9] independent observations for the PO [MA] auction. In each period, all subjects drew a new efficiency parameter.[1] Tables 1 and 2 summarize our experimental design.

Table 1   Experimental Design

Treatment	# Groups	# Subjects per group
Price-only auction	6	4
Multiattribute auction	9	4

Table 2   Experimental Design – Organization of the Treatments

Part	Competition?	Environment	Economy	# Rounds
1	No	Deterministic	$P\{\varepsilon = 0\} = 1$	5
2	No	Stochastic	$P\{\varepsilon = -2\} = P\{\varepsilon = 0\} = P\{\varepsilon = 2\} = 1/3$	5
3	Yes	Deterministic	$P\{\varepsilon = 0\} = 1$	10
4	Yes	Stochastic	$P\{\varepsilon = -2\} = P\{\varepsilon = 0\} = P\{\varepsilon = 2\} = 1/3$	10

We paid subjects a lump sum transfer of €5 for showing up and an additional reward equivalent to their earnings during the auctions.[2] In parts 1 and 2 [3 and 4], we exchanged points into cash according to the exchange rate 200 points = €1 [20 points = €1]. If the subjects had played according to the equilibrium strategies, they would have earned €16.40 on average including the show-up fee in both auctions. It turned out that in the PO [MA] auction, subjects earned €17.04 [€14.42] on average in approximately 2 [2.5] hours including the show-up fee. In order to help subjects better understand the environment, we programmed calculators that allowed subjects to calculate their profits conditional on winning and effort

---

[1] For the sake of comparability of the results, we kept draws constant across treatments.

[2] Paying every period as we did induces behavior towards risk neutrality. Paying according to one randomly selected period, instead, may increase subjects' willingness to take risks (Davis and Holt 1993).

level. Moreover, we showed them a visual aid that depicted their marginal costs of effort. After the subjects finished reading the instructions, they had to answer test questions so that we could be sure that they understood the rules of the game.

4. Results: Aggregate Data

Obviously, parts 3 and 4 are the most relevant parts to study the relative performance of the two auctions. Table 3 below presents the aggregate outcomes of the mechanisms in terms of total surplus, the buyer's utility, and the suppliers' profit. Table 4 includes results of Mann-Whitney U tests to establish whether differences between mechanisms are statistically significant.

Table 3   Aggregate Results

	Environment	**PO**		**MA**	
		Theory	Observed	Theory	Observed
	Deterministic	82.50	79.37	82.50	79.63
Total surplus	Stochastic	78.40	76.05	78.40	75.65
	Total	80.45	77.71	80.45	77.64
	Deterministic	66.25	62.98	66.25	69.51
Utility buyer	Stochastic	64.20	59.31	64.20	62.92
	Total	65.23	60.98	65.23	66.22
	Deterministic	16.25	16.4	16.25	10.12
Utility suppliers	Stochastic	14.20	17.06	14.20	12.73
	Total	15.23	16.73	15.23	11.42

The following conclusions emerge if we consider differences statistically significant at the 10%-level. Let us first consider efficiency. The two auctions are equally efficient. Both mechanisms generate a simi-

lar level of total surplus: 77.71 [77.64] for the PO [MA] auction. The experiment hence confirms the theoretical result from section 2.

CONCLUSION 1. *Both auctions are equally efficient.*

Table 4   Results of Mann-Whitney U tests

	Alternative hypothesis	Deterministic	Stochastic	Total
Total surplus	MA ≠ PO	p = 0.953	p = 0.637	p = 0.768
Utility buyer	MA > PO	p = 0.056	p = 0.131	p = 0.070
Utility suppliers	PO > MA	p = 0.034	p = 0.044	p = 0.014

What about the utility of the buyer? The MA auction dominates the PO auction in terms of buyer's utility. This result is not in line with the theoretical outcome of utility equivalence. The level of utility for the buyer in the MA auction in the experiment exceeds the theoretical outcome. This is due to the results in the deterministic part: 69.51 [66.25] in the experiment [theory]. The level of utility for the buyer in the PO auction in the experiment, on the other hand, lies below the theoretical level.

CONCLUSION 2. *The MA auction yields higher expected utility for the buyer than the PO auction.*

Finally, we consider the suppliers' profits. Since the total pie is the same for both auctions, the results of the experiment in terms of profits for the suppliers in both mechanisms are the mirror image of those for the buyers' utility: the PO auction dominates the MA auction. The difference is statistically significant both in the deterministic and stochastic environments. It applies for the PO auction that the results of the experiment exceed the theoretical results whereas this is not the case for the MA auction.

CONCLUSION 3. *The PO auction generates higher expected supplier profit than the MA auction.*

In order to study the robustness of conclusions 2 and 3 with respect to learning effects, we estimate the following random effects panel data model:

$$U_{jr}^{AE} = \alpha + \beta^{stoch}(E-1) + \beta^{MA}(A-1) + \gamma^{PO}r(2-A) + \gamma^{MA}r(A-1) + \delta_{jr}^{AE} + u_g^{AE},$$ (2)

$j = 1,2,\ldots,n^A$, $r = 1,\ldots,10$, $A = 1,2$, $E = 1,2$, where n^A is the number of triples that participate in auction A. In (2), U denotes the buyer's or supplier's utility, g the matching group, $j \in g$ the relevant group of 2 subjects, r the round, A the auction (where $A = 1$ for the PO auction, and $A = 2$ for MA auction), and E the environment (where $E = 1$ for the deterministic environment, and $E = 2$ for the stochastic one). The error terms δ and u are *iid* according to a normal distribution with zero mean, whereby u captures the panel structure of the data. The β parameters measure the effects of the environment and the auction type, while the γ parameters correct for potential time trends. Table 5 contains the parameter estimates for both the buyer and the seller.

Table 5   Parameter estimates for (2)

	Utility buyer	Utility winning supplier
Constant	60.054**	17.508**
Stochastic	−5. 547**	1.834
MA	8.625*	−9.407**
Time trend PO	0.673*	−0.308
Time trend MA	0.057	0.437

Note: ** and * denote statistical significance at the 1% and 5% level respectively.

Table 5 confirms that the MA auction is more [less] profitable for the buyer [suppliers] than the PO auction, but the time trends also suggest that this difference may vanish over time. Therefore, we re-ran the regression in (2), but using rounds before the end instead of rounds (so, round 1 is relabeled as round 10, round 2 as round 9, and so forth). Indeed, the difference between the auctions in terms of the buyer's and suppliers' utilities ceases to be statistically significant (see table 6).

Table 6   Parameter estimates for (2) where rounds are relabeled as rounds before the end

	Utility buyer	Utility winning supplier
Constant	67.459**	14.115**
Stochastic	−5. 547**	1.834
MA	1.842	−1.203
Time trend PO	−0.673*	0.308
Time trend MA	−0.057	−0.437

Note: ** and * denote statistical significance at the 1% and 5% level respectively.

CONCLUSION 4. *The differences between the PO auction and the MA auction in terms of the buyer's and the suppliers' profit vanish over time.*

5. Results: Individual Data

In this section, we will further explain the above aggregate observations on the basis of individual subjects' behavior. We begin with the subjects' effort choices. The subjects could choose a level of effort between 0 and 30. Figure 1 depicts the absolute difference between observed effort and the optimal effort level in parts 3 and 4. The data reveal that subjects stay close to the dominant strategy in both parts in each mechanism: The average absolute difference between the profit maximizing effort and the actual effort is always less than 2, apart from round 3 in part 3 of the MA auction. There are two reasons why the actual effort levels deviate somewhat from the optimal ones. First, subjects could only enter integers,

while the optimal effort levels are typically not an integer. Second, profits were rounded to the nearest integer. Because profit as a function of effort is flat at the optimum, effort levels one or two units above or below the optimal one could still generate (almost) the optimal profit level. So, we can safely conclude that subjects understood very well which effort level maximizes profits. In other words, differences between the two mechanisms are mainly explained by the subjects' bidding behavior.

Figure 1   Average Absolute Difference between Optimal and Observed Winner's Effort
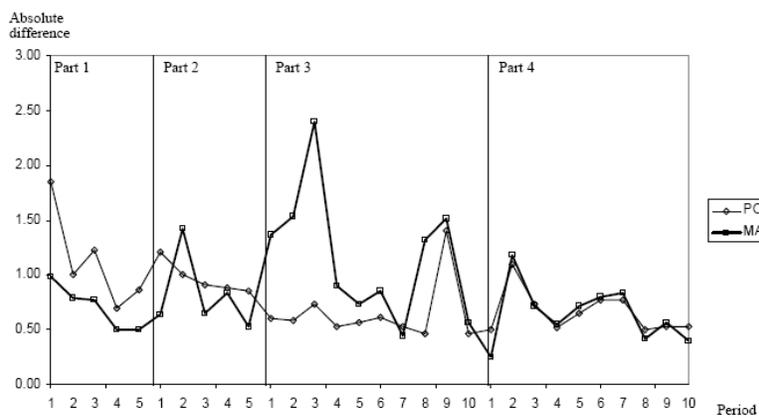
Figure 2 reveals that average (monetary) bids are close to equilibrium levels for the MA auction and somewhat lower than what the theory predicts for the PO auction. The latter observation contrasts with the overbidding that is usually observed in experimental first-price sealed-bid auctions (Kagel 1995). There are no noticeable differences between the deterministic and stochastic environments. When players are risk averse we expect to find more cautious bidding in the stochastic environment. In the PO auction, lower bids indicates more cautious bidding, whereas in the MA auction, the opposite is the case. Observe that in both mechanisms, there is only a slight hint of less aggressive bidding, because the grey line lies somewhat below [above] the bold black one in the PO [MA] auction. The differences are small though. Figure 2 also shows quite some variation in the bids in the sense that suppliers with the same efficiency level submit different bids. Therefore, it is unlikely that always the more efficient supplier wins, in contrast to what we find in theory. Indeed, table 7 indicates that this is the case.

Figure 2   (Monetary) Bids in PO and MA
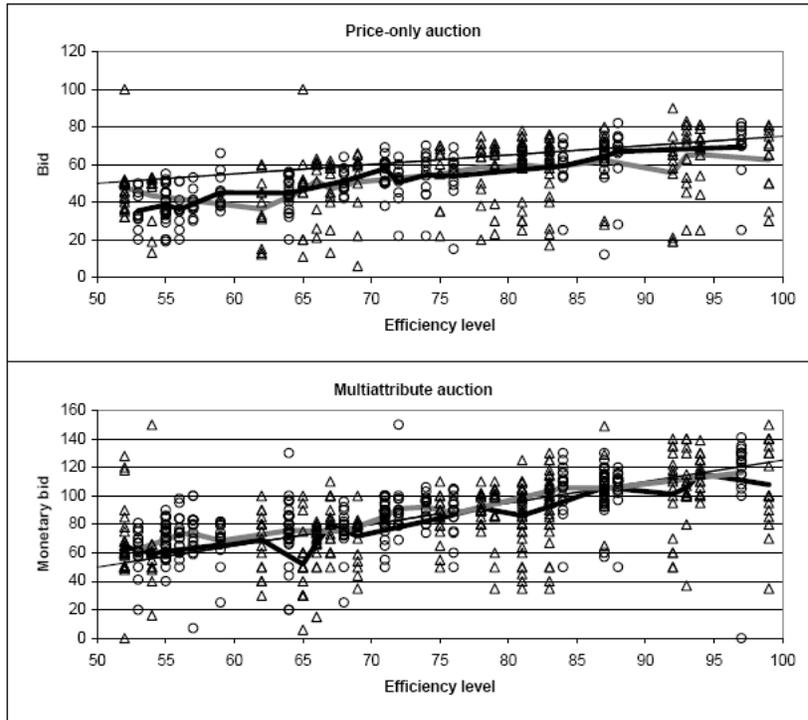
Notes: Open triangles [circles] indicate observed bids in the deterministic [stochastic] environment. The bold black [grey] line connects average bids conditional on efficiency type in the deterministic [stochastic] environment. The thin line depicts the risk neutral Nash equilibrium bids.

Table 7   Fraction More Efficient Supplier Wins

Mechanism	Deterministic	Stochastic	Total	Theory
PO	74.2%	80.0%	77.1%	100%
MA	81.7%	76.7%	79.2%	100%

The above observations explain most of our findings in the previous section. First of all, we noted that the two auctions are equally efficient. Second, in both auctions, total surplus is close to the theory. In section 2, we established that full efficiency will emerge if the buyer always selects the most efficient supplier, and let her exert effort according to (1). Because the winning supplier's effort is close to optimal, both

auctions are equally efficient if they are equally likely to pick the most efficient supplier. Table 7 shows that this is roughly the case, although the PO auction is somewhat more likely to select the most efficient supplier. Indeed, both mechanisms are very efficient despite the fact that only in about 75% of the cases the most efficient bidder wins. This indicates that the mechanisms are robust in the sense that the surplus losses are relatively small when the buyer does not award the contact to the most efficient supplier.

While the two auctions are equally efficient in the sense that the pie to be divided among the supplier and the buyers is about the same, the suppliers initially obtain a larger share in the PO auction than in the MA auction. The main reason is the following. Figure 2 indicates that average bids in the MA auction are at the equilibrium level, while those in the PO auction are somewhat lower than the equilibrium level. Proposition 3 shows that both auctions would have resulted in the same expected profit for the suppliers if they had played according to equilibrium. Because the winning suppliers in the PO auction pay their bids to the buyer, a below equilibrium bid implies that the suppliers' utility is higher than in equilibrium.

6. Conclusion

In this note, we have examined the relative performance of two procurement mechanisms as a means to allocate incentive contracts: price-only auctions and scoring auctions. We have done so both in theory and in the lab in an environment that differs in one important dimension from earlier laboratory experiments: In contrast to the latter, we gave bidders incentives to exert effort after the PO auction. In this sense, our experimental setting enables us to study an important aspect of procurement practice. Our main theoretical results are the following. The PO auction and the MA auction generate the same utility for the buyer, the same profit for the suppliers, and the same total surplus. Our experimental results show that PO auctions may not perform as poorly relative to MA auctions as suggested by previous experiments: The MA auction only initially outperforms the PO auction in terms of buyer utility, while total surplus is the same for the two auctions, and suppliers' profits are initially higher in the PO auction. We find that these initial deviations from theory disappear when subjects have had time to learn. These findings challenge the

European Commission's advice that EMAT should be the preferred format to "ensure that enterprises bear some responsibility for quality and performance of the completed works".

Future research may further explore the relative performance of MA and PO auctions. We have simplified the environment in three important ways. First of all, we have considered one-dimensional effort. It will be interesting to also explore settings in which the buyer can expend effort and/or submit bids on several dimensions, including cost reduction, delivery time, quality, and quantity. Second, the parameters of our environment are such that the MA auction that implements the socially optimal mechanism has a linear scoring function which is equal to the buyer's utility function. Che (1999) shows that in general, the buyer optimal MA auction underrewards effort relative to the buyer's utility function and more so for low levels of effort. A natural question is how simple, linear MA auctions perform relative to the buyer optimal MA auctions. This question is particularly interesting in the light of Onderstal's (2009) theoretical finding that in a broad class of settings, there is a linear MA auction that performs almost as well as the optimal auction if the number of bidders is high. Finally, we assume that output is ex ante contractible. However, in many practical settings, the tendered contracts are incomplete in that an a priori specification of the non-price dimensions is not feasible. The research of auctioning incentive contracts with incomplete specification of some dimensions will be a challenging avenue that could contribute to the unification of auction theory and the literature on incomplete contracts.

References

Asker, J., E. Cantillon. 2008. Properties of scoring auctions. *RAND J. Econom.* **39** 69-85.

Bajari, B., and G. Lewis (2009). Procurement contracting with time incentives: Theory and evidence. NBER working paper 14855.

Bichler, M. 2000. An experimental analysis of multi-attribute auctions. *Dec. Support Systems* **29** 249-268.

Branco, F. 1997. The design of multidimensional auctions. *RAND J. Econom.* **28** 63-81.

Burmeister, B., T. Ihde, T. Kittsteiner, B. Moldovanu, J. Nikutta. 2002. A practical approach to multi-attribute auctions. *Proceedings of the DEXA Workshops*, 670-674.

Che, Y.-K. 1993. Design competition through multidimensional auctions. *RAND J. Econom.* **24** 668-680.

Chen-Ritzo, C.-H., T.P. Harrison, A.M. Kwasnica, D.J. Thomas. 2005. Better, faster, cheaper: An experimental analysis of a multiattribute reverse auction mechanism with restricted information feedback. *Manage. Sci.* **51** 1753-1762.

Davis, D.D., C.A. Holt. 1993. *Experimental Economics*. Princeton, NJ: Princeton University Press.

D'Alpaos, C., M. Moretto, P. Valbonesi, S. Vergalli. 2009. "It is never too late": Optimal penalty for investment delay in public procurement contracts. FEEM working paper 78.2009.

Engelbrecht-Wiggans, R., E. Haruvy, E. Katok. 2007. A comparison of buyer-determined and price-based multiattribute mechanisms. *Marketing Sci.* **26** 629-641.

European Commission 2003. *Report and Recommendations of the EMAT Task Group: A methodology that permits contract award to the Economically Most Advantageous Tender.* Brussels, Belgium.

European Commission 2004. *A Report on the Functioning of Public Procurement Markets in the EU: Benefits from the Application of EU Directives and Challenges for the Future*. Brussels, Belgium.

Haruvy, E., E. Katok (2007). An experimental investigation of buyer determined procurement auctions. Working paper.

Herbsman, Z., W.T. Chen, W. Epstein. 1995. Time is money: Innovative contracting methods in highway construction. *J. Construction Eng. Manage.* **121** 273-281.

Kagel, J.H. 1995. Auctions: A survey of experimental research. J.H. Kagel, A.E. Roth, eds. *The Handbook of Experimental Economics*. Princeton University Press, Princeton, NJ.

Laffont, J.-J., J. Tirole. 1987. Auctioning incentive contracts. *J. Political Econom.* **95** 921-937.

McAfee, R. P., J. McMillan. 1986. Bidding for contracts: A principal-agent analysis. *RAND J. Econom.* **17** 326-338.

McAfee, R. P., J. McMillan. 1987. Competition for agency contracts. *RAND J. Econom.* **18** 296-307.

Onderstal, S. 2009. Bidding for the unemployed: An application of mechanism design to welfare-to-work programs. *Europ. Econ. Rev.* **53** 715-722.

Appendix

In this appendix, we provide proofs of propositions.

Proof of Proposition 1. If a supplier with efficiency level t wins the project, it is easy to check that her optimal effort is $e^{PO}(t) = ct\tau_p$. Her expected profit in the second period equals $ct(\tau_p)^2/2$ so that the profits are distributed according to a uniform distribution on the interval $[ca(\tau_f)^2/2, cb(\tau_f)^2/2]$. Because a supplier's profit represents the value of winning the auction, it is straightforward to check that B^{PO} constitutes an equilibrium strategy. Maskin and Riley (2003) show that this equilibrium is unique. It is readily verified that the PO auction implements the first-best optimal mechanism if $\tau_p = s$, so that, *a fortiori*, it is a socially optimal mechanism.

Proof of Proposition 2. For a given score σ, supplier i chooses effort e and transfer T solving

$$\max_{e,T} T - C(e, t_i) = T - e^2/(2ct_i)$$

s.t. $\sigma(e, T) = \tau_m e - T = \sigma$.

Substituting $T = \tau_m e - s$, the problem becomes

$$\max_e \tau_m e - \sigma - (e^2)/(2t_i).$$

Observe that $e^{MA}(t_i) = c\tau_m t_i$ is a solution. Note that this solution does not depend on σ or T. If supplier i wins, her utility equals

$$T - C(e^{MA}(t_i), t_i) = \tau_m e^{MA}(t_i) - \sigma - C(e^{MA}(t_i), t_i) = ct(\tau_m)^2/2 - \sigma.$$

Suppose that in equilibrium, $\sigma^*(t) \equiv \sigma(e^{MA}(t), T^{MA}(t))$ is increasing in t. If a supplier acts as if having signal t' instead of her true signal t, her expected utility is

$$U^{MA}(t,t') = \left[\frac{t'-a}{b-a}\right]^{n-1} \left[\frac{c(\tau_m)^2 t}{2} - \sigma^*(t')\right],$$

where the first term in the product on the RHS refers to the winning probability, and the second to the expected utility conditional on winning. The first-order condition of the equilibrium is

$$\left.\frac{\partial U^{MA}(t,t')}{\partial t'}\right|_{t'=t} = 0$$

which implies

$$(n-1)\left[\frac{c(\tau_m)^2 t}{2} - \sigma*(t)\right] - \sigma*'(t)(t-a) = 0.$$

It is readily verified that

$$\sigma*(t) = \frac{c(\tau_m)^2}{2}\left[t - \frac{t-a}{n}\right]$$

is a solution, which implies that

$$T^{MA}(t) = \tau_m e^{MA}(t) - \sigma*(t) = \frac{c(\tau_m)^2}{2}\left[t + \frac{t-a}{n}\right].$$

Observe that $\sigma*(t) \geq 0$ for all t, and that σ is strictly increasing in t, so that the outcome of the MA auction is efficient. Analogous to the proof of Proposition 1, it is readily verified that the MA auction implements the socially optimal mechanism.

Proof of Proposition 3. Let $s = \tau_p = \tau_m$. The buyer's expected utility in PO (U_B^{PO}) and MA (U_B^{MA}) equals

$$U_B^{FP} = U_B^{SA} = E\{B^{FP}(t^{[1]})\} = \frac{cs^2}{2}E\left\{t^{[1]} - \frac{t^{[1]}-a}{n}\right\} = \frac{cs^2}{2(n+1)}((n-1)b + 2a),$$

where $t^{[1]}$ is the highest out of n types.

References

Maskin, E., J. Riley. 2003. Uniqueness of equilibrium in sealed high-bid auctions. *Games Econom. Behav.* **45** 395-409.

Electronic Companion: Instructions

The original instructions were in Dutch. The instructions were the same for the price-only auction (PO) and the multiattribute auction (MA) unless we indicate otherwise.[1]

Welcome to this experiment on decision-making! You can make money in this experiment. Your choices and the choices of the other participants will determine how much money you will make. Read the instructions carefully. There is paper and a pen on your table. You can use these during the experiment. Before the experiment starts, we will hand out a summary of the instructions.

The Experiment

The experiment consists of four parts. In each part, you will earn points. At the end of the experiment, your points will be exchanged in euros. You will start with a starting capital of 5 euros.

During the experiment, you will compete with another participant to win a project. If you win the project, you will decide how much effort you will exert. The more effort you exert, the more points you will earn. However, effort is also costly.

The Costs of Effort

The costs of effort depend on how efficient you are. Your efficiency level is an integer between 50 and 100, where each number between 50 and 100 is equally likely. The higher your efficiency level, the lower your costs. The figure below presents the additional costs of effort if your efficiency level equals 75.

Additionele kosten

Effic. = 50
Effic. = 75
Effic. = 100
uw Effic.

Inspanning: -
Add kosten: -

Inspanning
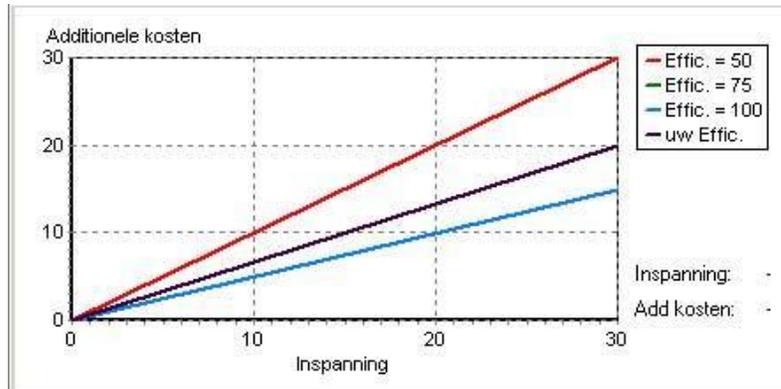
This figure indicates how much extra cost you will make if you provide one additional unit of effort. Observe that your costs are lower than at an efficiency level of 50, but higher than if your efficiency level had been 100.

An example illustrates the above. Suppose your efficiency level equals 50, and you deliver 3 units of effort. Your total costs are than $1 + 2 + 3 = 6$. The reason is that for the first unit of effort, the costs are 1, for the second unit 2, and for the third unit 3. If your efficiency level is 100 then your costs at the same level of effort are lower: $0.5 + 1 + 1.5 = 3$. Your costs at efficiency level 75 are somewhere in between.

To help you make these calculations, during the experiment, at the bottom-right corner of your screen, a calculator will appear.

The Benefits from Effort

Every unit of effort produces one unit of output, for which you get points. Suppose, for instance, that you receive 10 points per unit of output, and that you deliver 12 units of output. In that case you earn $10*12 = 120$ points.

Instructions Part 1

The first part of the experiment consists of 5 rounds. For every 200 points that you earn in this part, you will receive € 1 at the end of the experiment. In part 1, you won't yet compete with another participant for the project. In other words: you are guaranteed to execute the project. In each round, you must decide how much effort to expend in the project.

In part 1, you obtain 10 points for each unit of effort. Your profit from a project is the number of points you get for your effort minus your effort costs. As said, your costs depend on your efficiency level. At the bottom-right corner of your screen, a calculator will appear which calculates revenue, costs, and your profit if you type in specific effort levels. You can run as many calculations as you like.

Your efficiency level varies from one round to the next. Each round, the computer draws your new efficiency level, i.e., an integer between 50 and 100. Each number between 50 and 100 is equally likely.

Before we start part 1, let us return to the figure. This figure allows you to derive which effort level maximizes your profit.
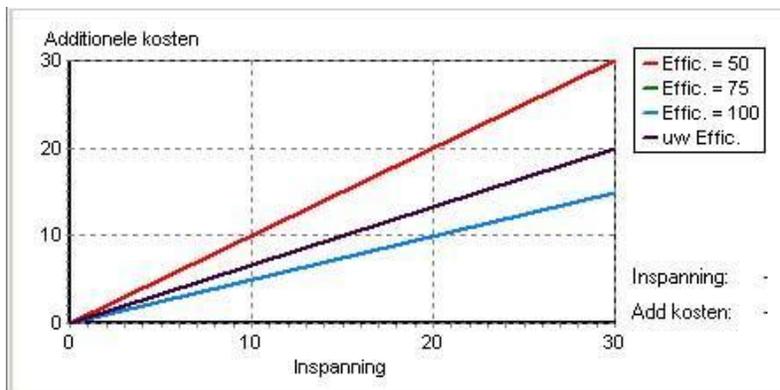
Suppose your efficiency level equals 50, and that you obtain 10 points for each unit of effort. Your profit is maximized if you choose an effort level of 10. Why? Let us start at effort level 0. It is more interesting to exert 1 unit of effort, because the costs for doing so are 1, while your gains from this unit of effort equal 10. The same holds true for the second unit of effort. This costs you 2 points, but its revenue is, once again, equal to 10. The same reasoning holds true for the third up to the tenth unit of effort: each unit yields 10, but costs less than that. Is it interesting to deliver an 11th unit? No: this will cost you 11 points, but generates only 10. The same holds true for any effort level above 10. Summarizing: it is optimal to expend 10 units of effort.

Instructions Part 2

The second part of the experiment is almost identical to the first. Also this part consists of 5 rounds. For every 200 points that you earn in this part, you will receive € 1 at the end of the experiment. The points

that you obtained in the first part will be added to your starting capital. The sum will be your starting capital for part 2.

Like in part 1, you exert effort in a project. This effort generates output. The main difference with part 1 is that in part 2, your output depends on the state of the economy, on which you have no influence. If the state of the economy is "neutral", then each unit of effort yields one unit of output. If the state of the economy is "good", then your output equals your effort plus two additional units of output. However, if the state of the economy is "bad", you lose two units of output compared to the neutral state. In short: you can determine your effort level, but you only know afterwards how much output it generates.

Let us give an example for clarification. Suppose you obtain 10 points per unit of output, and that you deliver 12 units of effort. In a neutral state of the economy, your output is 12 as well. In that case, you earn 10*12 = 120 points. If the economy is good, your output equals 12+2 = 14, so that you score 10*14 = 140 points. In the case of a bad economy, your output equals 12-2 = 10, which gives you 10*10 = 100 points.

Your efficiency level varies from one round to the next. Each round, the computer draws your new efficiency level, i.e., an integer between 50 and 100. Each number between 50 and 100 is equally likely.

Instructions Part 3

The third part of the experiment lasts for 10 rounds. For every 200 points that you earn in this part, you will receive € 10 at the end of the experiment (so 10 times as much as in the previous parts). Your earnings of the first two parts will be transferred to your starting capital of part 3. In this part, you compete with another participant for a project. Each round consists of two periods.

[PO: In the first period, you submit a bid. The same holds true for the other participant. The highest bidder will execute the project. He or she has to pas his or her bid, that will be subtracted from his or her current capital. If bids are identical, you both have 50% probability of winning the project. In the second period, the winner chooses how much effort to expend in the project. For each unit of effort, he or she

obtains 10 points. In this part, the state of the economy is always neutral. The output of the project is therefore equal to the winner's effort. The loser does not do anything in the second period.]

[MA: In the first period, you bid both an amount of money and an effort level. The same holds true for the other participant. These bids result in a score, which is equal to ten times the effort level minus the amount of money. The bidder with the highest score wins, unless the score is negative. If scores are identical, you both have 50% probability of winning the project. If you both offer a negative score, none of you will complete the project. The one who executes the project, exerts the effort promised in his or her bid in the second period. He or she receives the amount of money that he or she bid. The player with the lower score obtains zero points.]

In every round, you obtain a new efficiency level, in a same manner as in part 1. The same holds true for your competitor. Therefore, it is very likely that your competitor has a different efficiency level than you.

Finally, in each round, the computer selects at random the person with whom you will compete.

Instructions Part 4

The fourth part of the experiment also lasts for 10 rounds. For every 200 points that you earn in this part, you will receive € 10 at the end of the experiment (the same as in the previous part). The points that you obtained in the previous parts will be added to your starting capital. The sum will be your starting capital for part 4.

This part is similar to the previous one, the main difference being that the state of the economy can be good, neutral or bad. Each state is equally likely.

[MA: If you have the higher score, and if the state of the economy is good, you obtain an additional reward of 20, because your output is higher than promised. In the case of a bad economy, you get a fine equal to 20 points if you win the project.]

[PO: If the state of the economy is "neutral", then each unit of effort yields on unit of output. If the state of the economy is "good", then your output equals your effort plus two additional units of output. However, if the state of the economy is "bad", you lose two units of output compared to the neutral state.]