University of Cyprus
Department of Economics

Working Paper 04-2024

Feedback and Competition in Procurement e-Auctions

Niklas Klarnskou, Philippos Louis and Wouter Passtoors

Department of Economics, University of Cyprus, P.O. Box 20537, 1678 Nicosia, Cyprus
www.ucy.ac.cy/econ

Feedback and Competition in Procurement e-Auctions.

By Niklas Klarnskou[*], Philippos Louis[†] and Wouter Passtoors[‡]

Abstract:

We use a lab experiment to examine the effect of feedback on bidder behavior in procurement e-auctions. We compare 'Rank-only' to 'show lead bid' feedback, two regimes applied frequently by procurement professionals. A choice among the two is often based on rules-of-thumb that rely on initial 'bid compression', i.e. the spread of the bids submitted pre-auction. The use of such criteria finds no support in existing economic theory. A common assumption in theoretical auction models is that bidders face no opportunity cost from participating in a dynamic auction. This may not hold in situations where the expected value of a contract does not justify a long-time commitment to the bidding process on the part of bidders. In our experiment participants face the choice of remaining active in an auction vs. exiting and being rewarded with a diminishing outside option. Showing the lead bid accelerated bidders' learning. In the presence of opportunity costs, this can lead to substantially different outcomes conditional on the initial bid compression. With low bid compression, the bidder with the lowest cost wins more frequently, enhancing efficiency, but faces reduced competition by the others, which hurts the buyer's potential outcome. The opposite is true when bid compression is high. Rank only feedback achieved similar overall levels of efficiency, with higher benefits for the buyer. Crucially, these outcomes are not as sensitive to initial bid compression as in the case of 'show lead bid' feedback. A discouragement effect emerging in the 'Show-lead-bid', but not in the 'Rank-only' regime can explain these results.

Keywords: *Procurement auctions, feedback, opportunity cost, competition, bid compression, discouragement effect, lab experiment.*

JEL codes: *C92, D44, D83*

Acknowledgement: *We would like to thank Jacob Goeree, Pasha Adreyanov, colleagues at the School of Economics and Management at UCY, and participants in the Oligo Workshop 2023, for helpful comments and suggestions. We would also like to thank Neofytos Constantinou for planting the seed that started this project.*

[*] *Klarnskou,* email: *Niklas@klarnskou.com*
[†] *University of Cyprus,* email: *louis.philippos@ucy.ac.cy*
[‡] *Moneyball CPH,* email: *wouter@moneyballcph.com / Tilburg University,* email: *w.a.a.passtoors@tilburguniversity.edu*

1. Introduction

Without catching many headlines, the use of reverse auctions has been growing for more than two decades within the area of procurement, be it public or private. Whether applied in rising economic periods or uncertain and inflationary times, e-Auctions have proven an effective tool for procurement to mitigate the volatility in a competitive market. Private companies are applying the auction framework broadly across multiple categories and for contracts of all sizes, from multi-million agreements to smaller, day-to-day transactions. In fact, e-Auctions are more likely applied to smaller contracts where the time saved in the negotiation phase on both the buyer's and the bidders' side is of a higher relative value.

For practitioners, employing auctions over traditional negotiations ensures efficient business allocation in any economic climate, provided there is sufficient competition and the auctioned business is commercially appealing. Reverse dynamic auctions represent the purest form of negotiation, as they guarantee the auctioneer the best possible outcome on the auction day. However, it's important to note that without a commitment to award the contract based on the e-auction's outcome, the likelihood of the buyer achieving a favorable result significantly diminishes. It is the commitment to award that creates a genuine incentive for bidders to present their best and final offer.

Using one of many available software solutions organizations can run auctions to select suppliers for their needs. Some are highly customizable, allowing buyers to choose different auction formats, set reserve prices, define expressive bidding etc. Of course, all this variety of design options brings about the problem of how to choose the right design. Understanding the impact of such choices on an e-Auction's outcome becomes even more relevant following the rise of virtual assistants and AI assisted decision-making technology. Larger organizations that are conducting many auctions, are able to improve auction design through trial and error, leveraging their own internally build-up data sets. Even so, any mistakes can incur substantial costs to a company due to suboptimal supplier selection and/or overspending. The current source of remedies is sharing "best practices" with other practitioners in industry events (often organized by the software providers) or informally. Ideally though, they should be guided by carefully thought-out and empirically validated theory.[1]

It is well established in auction theory that the benefits of getting the details of the design exactly right are overshadowed by those of increasing the number of bidders (Klemperer, 2002). Beyond theory, practitioners are also aware of this fact. It is therefore no surprise that when planning an auction, they invest substantial resources in attracting all potential bidders. Furthermore, a high number of bidders is necessary for enhanced competition, but far from sufficient: in dynamic auctions, the benefits of nominal competition (high number of participating bidders) are only achieved if this translates into actual competition (high bidding intensity).

Attracting enough bidders is not always possible in the fast pace of business and can also be further hindered by capacity and resource constraints on the suppliers side. Practitioners therefore apply specific design features to an auction in an effort to mitigate the lack of bidders and still deliver satisfactory outcomes. Thus, nominal competition does not only depend on the number of bidders but also on specific design choices such as the available feedback. This is our study's main focus.

It is common practice in procurement auctions not to inform participants about every other bidder's current bid. In one common feedback regime, the leading bid is shown to everyone. Other times, bidders only receive rank feedback: how their own bid ranks compared to others. Knowing that one's bid is highly ranked may encourage the bidder to bid aggressively, while the opposite may happen if their bid is ranked low. If the leading bid is shown along with rank feedback, these effects may change, depending on how far off the leading bid lies. Whether any of this intuition is true will depend on what bidders

[1] For discussions of such issues see Hartley et al. (2004), Hur et al. (2006), Bruno et al. (2012), and references therein.

know about their competitiveness. Information about this can be revealed during the pre-auction stages, such as through one or multiple (often binding) requests.[2]

Understanding the interplay between the feedback provided during an e-Auction and the information revealed in the pre-auction stages is this project's main research question. We want to understand how 'rank only' vs. 'lead bid' feedback affects the behaviour of bidders, competition during the auction, and ultimately the performance of the auction as part of the procurement process.[3] To highlight the role of feedback on bidders' beliefs regarding their competitiveness vis-à-vis the others, our experiment employs a symmetric affiliated private value (APV) model, very similar to the one used by Kagel, Harstad & Levin (1987).[4] Under these conditions, bidders that know their own cost and initial offer, but have no information about others' initial offers or rank, should believe that they have an equal chance of having the lowest cost as anyone else. The feedback they receive at the beginning and during the experiment allows them to update these beliefs. And it is these effects that we seek to investigate.

The two feedback regimes we study are commonly applied by procurement professionals in practice. To choose between the two they typically employ self devised rules-of-thumb based on their companies' own experience. These, often rely on initial 'bid compression', i.e. the spread of the bids submitted pre-auction.[5] Nevertheless, there is no support for the use of such criteria in existing economic theory. In particular, standard auction theory predicts no effect on an auction's outcomes from providing either type of feedback.

An important caveat to these theoretical predictions is that standard auction models assume that participating in an auction entails no opportunity cost for bidders. In real auctions, bidders must commit valuable time to participate. The dynamic nature of such auctions allows bidders to withdraw (or simply become inactive) if they believe that their chances of winning are slim, and the required time can be spent more productively elsewhere. This is expected to be the case for contracts that are smaller in value relative to the time compensation of the executives involved in the bidding process.[6] But, how bidder beliefs are formed and updated during an auction can depend critically on the type of feedback received, as well as the distribution of pre-auction bids. Thus, introducing such opportunity costs and the option to exit into an auction model, may substantially alter its predictions. Unfortunately, the analysis of such a dynamic model quickly becomes intractable.[7] We examine simplified models to gain some insight, and these also provide some benchmarks for our experimental results. Still, our main approach to the question in hand remains experimental. In the same spirit of keeping the analysis tractable, we focus on symmetric APV winner-takes-all auctions.

---

[2] In procurement jargon these pre-auction stages are known as "request for quote" (RFQ) or "request for price" (RFP). They follow an initial "request for information" (RFI) stage where the buyer asks potential suppliers for information on their organization and products or services. Based on this information the buyer selects suppliers to send an RFP. Some buyers provide feedback to suppliers after the 1st RFP and ask for a 2nd RFP, which typically results in higher bid compression. The buyer finally invites a subset of the suppliers that provided competitive prices in the RFP to participate in an e-auction.

[3] Our analysis focuses on information and learning and how these are affected by the feedback provided. It is of course possible that behavioral biases in bidding behavior also get differentially activated across feedback regimes. Perhaps the most relevant in our case is a quasi-endowment effect, where a bidder bids more aggressively when provisionally awarded the item (Heyman, Orhun and Ariely, 2004). For more discussion of this and other behavioral biases that affect auction behavior see for example Erhart, Ott and Abele (2015), Podwol & Schneider (2016), Offerman, Romagnoli & Ziegler (2022), Babaioff, Dobzinski & Oren (2022), and references therein.

[4] The symmetric APV model is a special case of the affiliated value model (Wilson, 1977; Milgrom and Weber, 1982)

[5] See Larsen (2021), also quoted in the last section of this article.

[6] While e-auctions are sometimes used in large multi-million dollar projects that catch headlines and attention, most e-auctions concern smaller projects. For example, out of the more than 500 e-auctions conducted by Maersk, a global leader in transportation and logistics, 47% had a spend size below \$100k, with an additional 23% below \$1m.

[7] English auctions are often studied theoretically by looking at an equivalent clock auction. We cannot make use of this method, since in a clock auction there are no ranks or lead bid to provide feedback on.

In the experiment we design, five bidders participate in a reverse English auction. Their costs and initial offers are generated randomly but remain correlated. We draw the initial offers in a way that generates substantial (random) variation in initial bid compression, allowing us to test for the effect it has on the auction outcomes. During the auction, bidders are allowed to exit the auction and collect a payoff that is diminishing with time. This feature emulates the possible effect of opportunity costs from participation. We apply two treatments, each one with a different feedback regime: 'rank only' (RO) or 'show lead bid' (SLB).

To obtain theoretical benchmarks to compare our experimental results to, we computationally solve a simplified theoretical version of our setup, in which bidders first decide whether to participate in the auction or exit, and if they do not exit they submit a sealed-bid in a 2^{nd}-price auction. While this ignores the role of dynamics, it still provides a useful reference that helps formulate testable hypotheses for our experiment.

We find that on average bidders in the experiment do not exit as early as predicted by theory. In general they stay active longer in the RO treatment and exit faster and at a higher rate in SLB. In terms of efficiency, overall it is somewhat lower in the experiment than our theoretical benchmark, but there are no significant differences across treatments. Where we do find significant differences is in bidders' profits. These are substantially higher in the SLB treatment. While the level differs substantially from the theoretical benchmark, the direction of the effect is the one expected. The positive effect of SLB feedback on bidders' profits is realized when bid compression is low. This is explained by an apparent *discouragement effect*, akin to the one found in all-pay auctions and contests: the emergence of a clear leader can scare off competing bidders, drying out the effective competition in the auction. The existence of opportunity costs from participating in the auction introduces an all-pay auction element to the environment.

The paper proceeds as follows. In what remains of this section we discuss some of the related literature. In the next section we present a simple theoretical model of a 2-bidder auction that we hope illustrates how the feedback regime interacts with bidders' opportunity cost to affect the auction's performance. It also shows how a static version of the discouragement effect can emerge in the SLB environment when bid compression is low. We then proceed to explain our experimental design and procedures. We derive theoretical benchmarks based on a static analysis, similar to that of the 2-bidder model in the preceding section. Based on these we derive testable hypotheses to guide the reader when going through our experimental results. These are presented in the next section. Finally, we end the paper with some discussion of our results and related conclusions. The appendix contains some supplemental material.

1.1 Related Literature

This paper is closely related to existing experimental research that studies feedback in procurement auctions (Chen-Ritzo, Harrison, Kwasnica and Thomas, 2005; Haruvy and Katok, 2013; Elmaghraby, Katok and Santamaria, 2012;) including auctions that allow for expressive bidding (Adomavicius, Gupta and Sanyal, 2014) and the design of such auctions more generally (see for example Elmaghraby, 2007 for a survey). An important innovation of our experimental design is the incorporation of opportunity costs from staying in the auction, along with the possibility of exiting.

Given this innovation in the experimental design, our work also relates to the question of entry into auctions and how it is affected by the auction format. Early theoretical work in this area has focused on the case where potential bidders do not observe their type before deciding whether or not to enter the auction (McAfee and McMillan, 1987; Levin and Smith, 1994; Engelbrecht-Wiggans, 1993; Smith and Levin, 1996; Pevnitskaya 2004; Li and Zheng, 2009). This assumption is also reflected in some of the experimental work in the area (Smith and Levin 2002; Engelbrecht-Wiggans and Katok, 2005; Reiley 2005; Palfrey and Pevnitskaya, 2008; Ertaç et al. 2011). Often, though, it makes more sense to consider the case where potential bidders know their type when deciding whether or not to participate in an auction and condition their decision on that private information. This has been considered in theoretical work (Green and Laffont 1984; Samuelson 1984; Campbell 1998; Miralles 2008; Tan and Yilankaya 2006; Menezes and Monteiro, 2008; Lu, 2009; Cao and Tian, 2007 & 2010) but has received less attention from experimenters (Aycinena and Rentschler, 2018; Aycinena, Bejarano and Rentschler, 2018). Our work is closer to this last branch of the literature, in the sense that bidders are informed about their type when making the participation decision.

Throughout our analysis and experiment we maintain the implicit assumption that the e-auction stage is preceded by one or more stages where bidders submit binding bids,[8] which we take to be their initial offers, and these are correlated with their private cost. Theoretical analysis of indicative bidding, qualifying auctions and other two-stage mechanisms seems to indicate that pre-auction bids can be correlated to bidders' private values when they induce some form of commitment or entry costs (Ye, 2007; Quint & Hendricks). Experimental results lend support to this finding, but also indicate that commitment is not critical and even non-binding indicative bids can reflect private values (Kagel et al. 2008). These results provide justification for our assumption. It should be noted that for an environment with a substantial common-value element and an insider, Boone and Goeree (2007) show that a two-stage qualifying auction, an equilibrium exists where bidders bid their unconditional expected value in the first stage. Nevertheless, experiments by Boone et al (2009) do not find strong support for such behavior, with bidders apparently being drawn to the "babbling" equilibrium where their non-binding qualifying bids carry no information.

We need to mention two more auction theory papers that look at questions that are in spirit very similar to ours, although in different environments.

Ashkenazi-Golan et al. (2023), looks at private value two-stage auctions where the n highest bidders in the first stage can improve their bid in the second stage. They ask whether it is optimal to provide information about first-stage bid ranks or even bids. Such information can have two effects: one the one hand it may induce bidders to bid more aggressively in the second stage, but on the other hand, expecting to receive such information, bidders may bid less aggressively in the first stage. They find that the first effect becomes dominant for a sufficiently large number of bidders. We look at a single-stage dynamic auction and information influences in-auction competition, but it has a heterogeneous effect on bidders, inducing more competition from some, while leading others to exit. The overall effect is related to initial bid compression.

Hernando-Veciana & Michelucci (2018) look at an ascending auction model with an incumbent and a common-value component. They show that 'rushes' may occur in equilibrium, where all bidders simultaneously drop their demand. This is reminiscent of exits in our setting. What triggers these rushes in their common-value setting is the information revealed by the incumbent dropping out, which is very different from the opportunity cost motivation we examine. They propose a two-stage auction design to restore efficiency.

Finally, the literature studying contests and all-pay auctions predicts that a discouragement effect can emerge in the presence of asymmetries between competing agents (Baye et al., 1993; Gradstein, 1995; Baik, 1994; Stein, 2002; Konrad and Kovenock, 2009) and even with homogeneous agents, under some conditions (Fang et al. 2020). Lab experiments have verified the existence of such an effect (Davis and Reilly, 1998; Anderson and Stafford, 2003), especially in dynamic environments (Fonseca, 2009; Deck and Sheremeta, 2012). The existence of opportunity costs adds an 'all-pay auction element' to our environment, but the asymmetries between players, i.e. bidders' cost differences, are masked when 'rank only' feedback is provided. It is when bidders can see the lead-bid and asymmetries are large (i.e. low bid compression) that we observe a discouragement effect emerging.

[8] In real scenarios, these bid are not always binding and bidders may amend their initial bid before the auction starts. Even in these cases, the non-binding price and initial bid are typically highly correlated.

2. Theoretical considerations

Our analysis departs from standard auction theory models in some aspects: bidders face an opportunity cost from participating in the auction and can choose to exit; initial offers are determined (exogenously) at a pre-auction stage and are correlated with bidders' cost; the feedback regime determines what bidders know about their competitivenes at the beginning of, and during the auction. In this section use a simple model to illustrate how these elements of the environment interact with each other.

Consider the following auction exit game.[9] Two risk-neutral bidders participate in a reverse second-price sealed-bid auction. In the first stage they decide whether to exit the auction or not ($e_i \in \{0,1\}$). If they exit ($e_i = 1$) they receive a fixed payoff $\bar{u} \in (0,1)$. In the second stage bidders are not informed about others' exit decisions. Any bidder that does not exit in the first stage submits a bid. Given the second-price nature of the auction bidders have a dominant strategy that is to submit a bid equal to their cost (Vickrey, 1961). The winner receives a payment equal to the second highest bid (i.e. the other bidder's cost) or the reservation price if the other bidder exits in the first stage, and pays their cost. The reserve price is set to $P = 2$ for convenience.[10] A bidder that does not exit but loses the auction receives zero payoff.

For simplicity we assume that bidders' costs take the values $c_L = 0$ or $c_H = 1$ and each of the two bidders has an ex ante equal probability of having the lowest cost. Let $\theta_i \in \{L, H\}$ represent a bidder's state, with $\theta_i = L$ in case bidder i has the lowest cost, $\theta_i \neq \theta_j$ and $\Pr(\theta_i = L, \theta_j = H) = \frac{1}{2}$. Since bidders submit a bid equal to their cost, if no bidder exits, the one with the lowest cost wins the auction. The bidder with the highest cost can only win if the other bidder exits in the first stage. These are simplifying assumptions that are done here only to keep the model simple and to focus on the first stage and the bidders' decision to exit or not.[11]

In the first stage bidders do not know whether their cost is the lowest. An *initial offer* $2 + \rho_i$ is generated for each one, where $\rho_i \in [0,1]$ is a random draw from a distribution F_{θ_i} conditional on the bidder's state θ_i, with pdf:

$$f_\theta(\rho) = \begin{cases} 2 - 2\rho, & \theta = L \\ 2\rho, & \theta = H \end{cases}$$

These are triangular distributions with the mode on the lower and upper bound for $\theta = L$ and $\theta = H$ respectively. Conveniently, after observing their initial offer and applying the Bayes rule, bidder i updates their belief about having the lowest cost to $\Pr(\theta_i = L|\rho_i) = 1 - \rho_i$.[12]

Bidders observe their initial offer privately. Furthermore, we consider two feedback regimes for the first stage. In *rank only* (RO), bidders are also informed of their initial offer's rank $r_i \in \{1,2\}$, indicating whether their initial offer is 1st (the lowest) or 2nd: $r_i = 1 \Leftrightarrow \rho_i < \rho_j$. In *show leading bid*

[9] We refer to auction "exit" instead of "entry", which the common nomenclature in this literature, just to stay in line with our experimental design, where the auction is dynamic and bidders can leave at any point in time. For the static environment studied in this section, choosing to enter or not exit the auction are equivalent.
[10] Setting the reserve price at $P = 2$, equalizes the profit from winning the auction with the highest cost if the bidder with the lowest cost exits ($P - c_H$), and that from winning with the lowest cost with no bidder exiting ($c_H - c_L$).
[11] In a reverse second-price auction bidding one's private cost (or progressively up to that in the strategically equivalent reverse English auction) is a dominant strategy. Hence, in such an auction bidders do not need to form any beliefs about others' costs and bidding strategies.
[12] We define the initial offer as $2 + \rho_i > P$ just to make it clear that bidders still need to submit a bid in the second stage in order to win the auction. All the relevant information that matters to the bidder is contained in ρ_i, which they can perfectly infer by observing the initial offer. We therefore use the two interchangeably.

(SLB) the bidder with the higher initial offer ($r_i = h$) is also informed about the other bidder's initial offer (the 'leading bid'). To simplify notation from now on we use the rank to indicate the bidder: $i = r_i$.

In this setup, maximum allocative efficiency is achieved if the bidder with the lowest cost does not exit the auction. As we discuss below, this is not always the case, as bidders only receive a noisy signal about the rank of their cost. It does not matter for allocative efficiency if the bidder with the higher cost exits. From the perspective of a spending minimizing buyer, exits are always bad, as they reduce competition and lead to a higher price paid to the winning bidder. Hence, the performance of auction under each feedback regime hinges upon the exit behavior induced by the information provided to the bidders.

Remark 1: As bidders are ex ante symmetric, the decision to exit depends solely on their belief about having the lowest cost. *A bidder exits unless their belief about having the lowest cost is high enough.*

Remark 2: Given the information structure and the rank information provided in both feedback regimes, bidder 1 is always more optimistic than bidder 2 about their chances of having the lowest cost. If bidder 2 is optimistic enough to stay even if bidder 1 does not exit, then bidder 1 will be even more optimistic and not exit. If bidder 2 is not optimistic enough and exits, then bidder 1 can win for sure and is therefore better off not exiting. Therefore, *bidder 1 never exits in the first stage*. This is straightforward in the simple two bidder model, but is also true in the model with multiple bidders we examine later on: higher ranked bidders are more optimistic and therefore if a bidder does not exit, neither will any bidder better ranked than them.

Given the above we focus attention on bidder 2. We can show the following:

- In RO, bidder 2 's optimal choice is:

$$e_2^{RO}(\rho_2) = \begin{cases} 1, & \bar{u} > \frac{1}{3} \\ 1, & \bar{u} \leq \frac{1}{3} \ and \ \rho_2 > \hat{\rho}(\bar{u}) \\ 0, & otherwise \end{cases}, \qquad \hat{\rho}(\bar{u}) = \frac{3\bar{u} - 1}{2\bar{u} - 1}$$

- In SLB, bidder 2 's optimal choice is:

$$e_2^{RO}(\rho_2, \rho_1) = \begin{cases} 1, & \bar{u} > \frac{1}{2} \\ 1, & \bar{u} \leq \frac{1}{2} \ and \ \rho_2 > \tilde{\rho}(\rho_1, \bar{u}) \\ 0, & otherwise \end{cases}, \qquad \tilde{\rho}(\rho_1, \bar{u}) = \frac{\rho_1(1 - \bar{u})}{\rho_1 + \bar{u} - 2\rho_1\bar{u}}$$

When $\bar{u} > \frac{1}{2}$, in both feedback regimes bidder 2 exits in the first stage irrespective of their initial offer. Hence the feedback regime has no effect on the outcome. We therefore focus on the more interesting case where $\bar{u} \leq \frac{1}{2}$.

An important concept for the analysis here and later in the paper is that of *bid compression*, i.e. how close initial offers are to each other. In this simple model one can think of it as the ratio between initial offers $\frac{\rho_1}{\rho_2}$. The closer the initial offers are to each other, the higher the ratio, or in other words, the higher the bid compression.

Remark 3: By construction, bid compression affects behavior in the SLB feedback regime, but not in the RO feedback regime. Notice that $\tilde{\rho}(\rho_1, \bar{u})$ is increasing in ρ_1 and, by definition, $\rho_1 < \rho_2$. Thus, the condition for bidder 2 to exit in SLB becomes more restrictive as the two initial offers are closer together, or, in other words, when initial *bid compression* is high. So, *if bid compression is high enough, a bidder with a specific initial offer may choose to exit in RO but not in SLB. On the other hand, with low bid compression, if a bidder with a specific initial offer does not exit in RO he will also stay in the auction in SLB*. Formally, we can show that:

- For $\bar{u} < \frac{1}{2}$:
  - If $\frac{\rho_1}{\rho_2} > \frac{1}{2}$ *then*: $e_2^{SLB}(\rho_2, \rho_1) = 1 \Rightarrow e_2^{RO}(\rho_2) = 1.$
  - If $\frac{\rho_1}{\rho_2} \leq \frac{1}{2}$ *then*: $e_2^{RO}(\rho_2) = 0 \Rightarrow e_2^{SLB}(\rho_2, \rho_1) = 0.$

The intuition behind the above result is the following. Being ranked 2nd does not necessarily mean that a bidder does not have the lowest cost, but it means that this is less likely to be the case. The likelihood decreases further with the bidder's initial offer. In RO, this is all the information available to the bidder when deciding whether or not to exit. In SLB, there is more information available. When the 2nd bidder observes the 1st bidder's initial offer to be much lower than their own, i.e. bid compression is low, the likelihood of having the lowest cost becomes even smaller. On the contrary, observing that the 1st bidder's initial offer is close to their own, i.e. bid compression is high, is good news for the 2nd bidder, as it means that they perhaps have the lowest cost but got an "unlucky" draw for their initial offer.

What we describe above is essentially the possibility for a discouragement effect arising. In particular we find that in RO the bidder ranked 2nd can be discouraged and exit by having a high initial offer. This decision will not depend on bid compression, as in RO a bidder knows nothing about the other bidder's initial offer. The opposite is true in SLB, where the discouragement effect is induced by the observation of the other bidder's initial offer: the further ahead the other bidder is, the stronger the effect.

In this simple two bidder model, bid compression is easily defined as the ratio between the two bidders' initial offers. Later when we examine a model with multiple bidders and a different cost distribution, and when looking at the experimental data, we will use the standard deviation of initial offers as a proxy for bid compression.

Remark 4: Bid compression affects allocative efficiency and buyer's spending under both feedback regimes. Given the previous remark, this is quite evident for SLB, but also holds in RO. Recall that both efficiency and spending depend on the probability of a mistake exit, i.e. the bidder with the lowest cost exiting the auction. In RO, the probability of the 2nd bidder exiting does not depend on the 1st bidder's initial offer, but if exit occurs, the probability of this being a mistake, conditional on the two initial offers being close, is high. The opposite is true when the two initial offers are far. Therefore, we expect to observe different efficiency and spending, even in RO, when we look separately at cases where bid compression is high vs. low.

The theoretical model presented in this section is meant to fix ideas and illustrate the strategic considerations present in our experimental design, how they are affected by the feedback regime, the role of bid compression, and how all of this can affect an auction's performance. In this simplified form the model does not lend itself for meaningful comparisons regarding performance across feedback regimes. We return to these issues after first introducing our experimental design. Based on that we set up and solve computationally a richer theoretical model that is closer to the auctions we conducted in the lab. We use that to provide benchmarks and hypotheses related to our experiment. All this is done in the following section.

3. The Experiment

3.1 Experimental design

The experiment took place online with 80 participants recruited from the UCY-LExEcon lab subject pool.[13] This is based at the University of Cyprus and participants are all students at the various departments of the university.[14] We conducted a total of eight separate sessions, with 10 participants per session. The experiment was programmed and run in zTree (Fischbacher, 2007), using zTree-Unleashed (Duch et al., 2020) and Zoom for the online implementation. An outline of the design is provided in

Treatment	Feedback	# of bidders per auction	# of auctions per round	# of rounds per session	# of sessions
SLB *(show lead bid)*	*Rank and lowest offer*	5	2	40	4
RO *(rank only)*	*Rank only*	5	2	40	4

. A translation of the instructions for the experiment provided to participants can be found in the Appendix. The original text in Greek is available upon request. We discuss further details of the design below.

In all sessions, subjects connected on to a Zoom meeting, through which they received instructions.[15]

Table 1: *Experimental Design*

The experimental task involved bidding in 40 auctions, one in each round of the experiment. All auctions are of the English reverse auction format, which is the most common format used in procurement e-auctions (Larsen, 2021). Subjects are randomly placed in groups of 5 bidders[16] at the beginning and the groups remain fixed in all rounds of the experiment.

Treatment	Feedback	# of bidders per auction	# of auctions per round	# of rounds per session	# of sessions
SLB *(show lead bid)*	*Rank and lowest offer*	5	2	40	4
RO *(rank only)*	*Rank only*	5	2	40	4

Each auction had a duration of 60 + x seconds, where x was a random integer between 1 and 10 drawn in each round and not revealed to the subjects. This ensured a random ending time for the auction and was used to discourage sniping.[17] Bidders submit offers for a virtual item in real time and as many times as they wish. Each bidder's new offer must be lower than her previous one but does not have to be the lowest one.

[13] Recruitment was done using ORSEE (Greiner, 2015).
[14] More information of the the UCY LExEcon lab and subject pool can be found at https://www.lexecon.ucy.ac.cy/
[15] Participants turned their cameras on upon entering the meeting individually, when their ID was checked. After that they switched off the camera, replaced their screen name with a randomly assigned number and moved into a "breakout room" while waiting for the remaining participants. This process assists in preserving anonymity.
[16] As this is a procurement setting, bidders are sellers/suppliers competing with their price offers to supply a good. Nevertheless, we will use the term 'bidders' which is more common in the study of auctions.
[17] See for example Roth & Ockenfels 2002. Real e-auctions usually extend the time limit if a bid is submitted near the end of the auction to achieve the same goal. This is not practical in the lab where multiple auctions are conducted simultaneously and repeatedly.

The bidder that has submitted the lowest offer within the time limit wins the auction. She is paid the offer amount minus her cost (see below). All other bidders remaining active by the end to the auction get a payoff of zero. At any point during the auction, bidders have the option to click on a button and exit the auction. Doing so results in a payoff which is decreasing with time. It starts at 30 points and decreases by 1 point in 2 second intervals. When a bidder exits, they become inactive for the remainder of the auction. Exit in any round, though, does not affect participation in any of the following rounds.

Bidders' costs are determined as follows. In each round the computer randomly picks 1000 consecutive integers between 2000 and 8000. Each possible such sequence is equally likely. From these integers one is picked, with a uniform probability, for each bidder to represent her cost. Thus, each bidder knows that the other bidders' cost is at most 999 points higher or lower than their own.[18]

At the start of every round each bidder is assigned an initial offer in the following way: the computer picks randomly and uniformly an integer between 1500 and 3500 for each bidder and adds it to their cost. 10 seconds before the auction clock starts, bidders are informed about their individual cost, their initial offer and their offer's rank (based on the initial offer), with the lowest offer being ranked 1st and the highest offer 5th. In the SLB treatment bidders are also shown the amount of the lowest offer. Notice that given the way initial offers are determined, a bidder whose initial offer is ranked higher than that of another one is more likely to have a lower cost than that other bidder. Overall, the correlation between initial rank and cost rank is .325.[19] It is important to note for the subsequent analysis that in the cases where the spread of the initial offers is low, the correlation remains positive, but is considerably lower than in the opposite case (.138 vs. .475). Of course, participants are not informed about the spread of initial offers.

In both treatments, once the auction starts bidders are shown real time information about their current offer's rank with respect to other active bidders. Whenever a bidder submits a new offer everyone's rank information is updated accordingly. In the SLB treatment bidders also observe the current leading bid, i.e. the amount of the lowest offer submitted by any active bidder in the auction. Bidders that exit do not receive real time information for the remainder of the auction. When a bidder exits, remaining bidders' ranks are updated based on their current offers of active bidders, i.e. the exiting bidder's offer is withdrawn.[20]

At the end of the time limit all bidders, whether active or inactive, are informed about the auction result. Specifically, they are shown the amount of the winning offer, the amount and rank of their own final offer, whether they won, lost or exited, and their payoff.

3.2 Theoretical benchmarks

Before we are ready to formulate testable hypotheses for our experimental results, it is useful to derive some theoretical benchmarks, specifically tailored to our experimental design. We do this here.

In the auction environment of our experiment, in each point in time subjects need to decide whether or not to stay in the auction and, if they do, whether or not and by how much to lower their current offer. All that while information about their own rank and (in the SLB treatment) the leading bid is continuously changing. This is not very different than what bidders in real procurement e-auction need to deal with. This realism comes at a cost. It makes the analysis of the strategic environment using equilibrium game theoretic tools intractable. Nevertheless, it is possible to take a "partial equilibrium"

[18] Kagel et al. (1987) employ a very similar method to induce affiliated values in their experiments.

[19] See also Table 5 in the Appendix.

[20] Updating all remaining bidders' rank when a bidder exits could convey to them some information about the remaining number of bidders, if the exiting bidder's current offer was ranked higher than another one. Note, in our theoretical analysis bidders are not informed about how many other bidders exit. Still, we choose to update ranks to avoid deception and stay closer to real-life implementations of such auctions.

approach to derive some benchmarks and formulate hypotheses to aid the interpretation of our experimental results.

3.2.1 No opportunity cost (i.e. no exit)

A first approach to understand what may determine bidders' behavior in this environment is to think of the simpler case where there is no opportunity cost from staying in the auction. In other words, one could consider the same environment, but without the possibility of exit. In such an environment, any rational bidder that is not ranked 1st should continue lowering their offer until reaching their cost. This is true irrespective of whether bidders receive feedback on the leading bid or not. As is well known, this standard reverse English auction is theoretically equivalent to a reverse 2nd-price sealed-bid auction: the winner is the bidder with the lowest cost and the price is equal to the 2nd lowest bidder's cost.

This result also implies that the auction outcome is always efficient: the bidder with the lowest cost wins the auction. In this scenario *without exit neither the feedback regime nor bid compression should affect the efficiency of the auction*. Similarly, they *should also not affect bidders' expected profit*.

Infact, the latter we can calculate exactly in our setup. Conditional on having the lowest cost, and hence winning the auction, a bidder's expected profit is then given by the expected difference between the 2nd lowest and lowest cost. As costs are uniformly distributed, this difference between order-statistics follows a Beta distribution with parameters α given by the difference in orders ($\alpha = 2 - 1$) and $\beta = n - \alpha + 1$, where n represents the number of bidders. For five bidders and a 1000 point interval this means the expected profit for the winning bidder is

$$E(\pi_{no\ exit}|winning) = \frac{1}{6} \cdot 1000 = 166.66 \text{ points.}$$

There are two things to note about this benchmark. The first relates to treatment differences and the second relates to how it may be affected by exits.

For the theoretical equivalence between the reverse English auction and the 2nd-price sealed bid auction to hold, it must be the case that bidders in the dynamic setting adjust their offers continuously by infinitesimal amounts. More realistically, it would require offer decrements to be small and very close to whatever is currently the leading bid. This is not very demanding in the SLB treatment, as bidders know exactly what the leading bid is and can adjust their offer accordingly to match it. On the other hand, in RO this is not possible. Given the time constraints it is reasonable for bidders to lower their offers in non-minimal decrements, which may result in undercutting the leading bid by a substantial amount. This would result in overall lower profits for bidders in the RO treatment. In their experiment, Elmaghraby et al. (2012) indeed find that prices tend to be lower with RO versus full-price feedback. They attribute this to jump bids as well as a similarity of the problem of bidders under RO to that presented in sealed-bid first-price auctions.

For the expected profit calculation above we assume 'no exit'. In fact, the calculation still holds if we only assume that the bidders with the lowest and 2nd lowest cost remain active in the auction. Of course, bidders do not know the rank of their cost and can only infer that by the information they receive at the start of the auction and the feedback they observe during the auction. If such inferences lead to the exit of bidders before their offer reaches their cost, expected profits may be different. This is explained in more detail below.

3.2.2 Static auction with opportunity cost (SAWOC)

To obtain some benchmarks of bidders' exit behavior we turn to a model where we dispense of all dynamics, but assume that bidders have an opportunity cost from participating in the auction. This is a variation of the simple two player model presented earlier that aligns more with the experimental setting. We now consider five risk-neutral bidders that participate in a reverse second-price sealed-bid auction. Bidders have affiliated costs determined in the same way as in our experimental design described above. Initial offers, ranks and the leading bid (in SLB) are also determined as in the experiment. In a first stage, after receiving information about their cost, initial offer, rank and, in SLB, the lowest initial offer, bidders must decide simultaneously whether or not to submit a new and final offer in the second stage, or exit and receive a fixed payment of 30 points.

In such a 2nd-price auction it is a dominant strategy for all bidders that do not exit to submit an offer equal to their cost. The outcome of the auction, who wins, the price, and expected profits, depends on bidders' exit decisions. In particular, the winner is the bidder with the lowest cost among the ones that do not exit. The price is the 2nd lowest cost among the ones that do not exit. If the two bidders with the overall lowest costs remain in the auction, the result is the same as in the case with no exit. But if either of the two decides to exit, results will differ.

First, notice that the option to exit may result in lower than optimal allocative efficiency if the bidders with the lowest cost decide to exit. Similarly, the winner's expected profit can also be substantially higher if either the bidder with the lowest cost exits, or she stays in the auction but the bidder with the second lowest cost exits. If the probability of any of this happening depends on the feedback received by bidders, both allocative efficiency and the winner's profits may differ between the two experimental treatments.

From the analysis of the model with no opportunity cost (i.e. no exit) we can see the following. Before receiving any information on their costs, initial offers and ranks, bidders are symmetric. Thus their ex ante expected profit from the auction is

$$E(\pi_{no\ exit}) = \frac{1}{5} \cdot E(\pi_{no\ exit}|winning) = 33.33 \text{ points.}$$

This means that, assuming risk neutrality, it is individually rational for bidders to participate ex ante. This is not the case anymore after bidders receive information about their cost, initial offers, and rank. Based on this information, bidders can update their belief about having the lowest cost, as they do in the simple two player – two type model we discussed. The bidders with the highest initial offers, as revealed by their rank and the leading bid (in SLB) will update their beliefs regarding their chances of winning downwards and may therefore choose to exit (see *Remark 1* and *2* for the simple model).

Let us first consider the RO treatment. Bidders observe their initial offer's rank before deciding whether or not to exit. Recall that initial offers are correlated with bidders' costs, and therefore a bidder ranked higher than another means that it is more likely for their cost to be lower than that of the other. Therefore, a bidder ranked 5th should adversely update their belief about the probability of winning the auction. In particular, given the parameters of the experiment the expected profit of such a bidder is:

$$E\left(\pi_{w\ exit}\middle|ranked\ 5^{th}\right) \approx 8.2 \text{ points.}$$ [21]

Therefore, assuming that all other bidders stay in the auction, it is optimal for the bidder ranked 5th to exit the auction and collect the fixed award of 30 points from doing so. Taking this into account we can calculate the expected profit of the bidder ranked 4th and obtain:

$$E\left(\pi_{w\ exit}\ \middle|\ ranked\ 4^{th}, 5^{th}\ has\ exited\right) \approx 25.1 < 30 \text{ points.}$$

Hence it is again optimal for the bidder ranked 4th to exit the auction. Given this, it turns out that the expected profit of the remaining bidders is:

[21] We obtain this and similar numbers using Monte Carlo simulations. The exact calculation of these figures while possibly tenable, is out of the scope of our paper.

$$E(\pi_{w\ exit} \,|\, ranked\ 3^{rd} or\ above,\ 4^{th} and\ 5^{th}\ have\ exited) > 30 \text{ points.}$$

To sum up, in a 2nd-price sealed-bid auction with exit with RO feedback, bidders with initial offers ranked 4th and 5th exit the auction and all other bidders submit an offer equal to their cost. The winner of the auction is the active bidder with the lowest cost and the price is equal to the 2nd lowest cost among the active bidders.

In the SLB treatment, observing the lowest initial offer provides more information to the bidders. They can therefore condition the decision to exit on their rank (as in RO) as well as on the leading bid information. So, unlike in RO, it is possible, for some initial offer profiles, to observe a bidder exiting ranked at any position up to the second, or the opposite: even the bidder ranked 5th to stay in the auction. Still, for any profile of initial offers, if a bidder exits, so will all bidders ranked worst, and if a bidder stays, so will all bidders ranked better.

We make use of Monte Carlo simulations to calculate a bidder's expected payoff conditional on their rank and the distance between their initial offer and the leading bid. Comparing that to the value of the outside option provides threshold values for the initial offer – leading bid distance, conditional on a bidder's rank.

From above we have computationally determined exit behavior in the SAWOC model for both the RO and the SLB feedback regime. We use this in a simulation exercise to obtain benchmarks regarding bidder's frequency of exit and the performance of the auction. These are presented in **Table 2**. The first two rows of the table show the frequency of exit under each feedback regime. The next two show the frequency of efficient outcomes, namely how often the winner of the auction was the bidder with the lowest cost. The last two rows show average winner's profits. The simulations also allow us to form some intuition on the role of bid compression. In the 2nd and 3rd columns in the table we split the sample equally between the cases with the lowest and highest bid compression. Here (and later for the experimental results) we measure bid compression by the standard deviation of initial offers.

Regarding exit behavior, the model predicts more frequent exits in SLB compared to RO. The intuition behind this is that the additional information provided to bidders in SLB allows worst bidders to more strongly update their beliefs about winning the auction. This means that, unlike the RO, even bidders that are ranked 3rd, or even 2nd sometimes choose to exit. This is more likely to happen when bid compression is low and therefore it may be that the bidder ranked 1st is more clearly ahead of the others. This is essentially a manifestation of the discouragement effect. As shown above, in RO the bidders ranked 4th and 5th always exit, and therefore bid compression plays no role here. These results are correspond to the similar finding in the simple model analyzed in Section 2 (Remark 3).

Table 2: *Simulated exit behavior and outcomes in the SAWOC model*

		Whole sample	**Low bid compr.**	**High bid compr.**
Exit	*RO*	.400	.400	.400
	SLB	.521	.638	.403
Efficient Outcomes	*RO*	.793	.867	.720
	SLB	.876	.852	.899
Winner's Profit	*RO*	222.5	213.9	231.1
	SLB	670.6	906.9	434.6

In terms of allocative efficiency, SLB feedback leads on average to slightly better results, although bid compression operates very differently in the two regimes. When bid compression is low, the ranking of initial offers is more closely correlated with the ranking of costs, and vice versa. Therefore, in RO, where it is always the worst two ranked bidders to exit, there are fewer "wrong exits" when bid compression is low. In SLB, with high bid compression bidders that are not well ranked can still observe the lead bid, which is likely to be close to their initial offer. Thus, even the bidders with the highest initial offers can see that the best offer is not much lower and therefore they maintain relatively optimistic beliefs about their chance of winning, and stay in. This leads to fewer "wrong exits" under high bid compression in SLB.

Recall that with no exit the expected winner's profit is about 167 points. In the SAWOC model we find that under RO it is somewhat higher, while under SLB it is substantially higher. This is explained by exits. Suppose the bidder with the lowest cost stays active (i.e. the outcome is efficient) but the one with the second lowest cost exits. This leads to a higher winner's profit. The higher exit rate in SLB leads to situations like this more frequently, in particular when bid compression is low.

We should emphasize that these benchmarks are obtained from a static model where feedback only affects the exit decision at the start of the auction. In our experiment, we employ a dynamic English auction with feedback, either RO or SLB, provided throughout the auction, and bidders being allowed to exit at any point. While the SOWAP model can help us build some intuition about how bidders can learn from feedback and behave in a dynamic auction, the experimental results can differ substantially from these benchmarks.

Next, we provide some hypotheses about bidders behavior in the experiment based on these theoretical benchmarks and our intuition. These should not be viewed as a way to validate a specific theoretical model, as any of our models differs substantially from the experimental setup. They serve mostly to organize, and later present, our results.

3.3 Hypotheses

Based on the intuition that motivated this work we expect exit behavior to differ across treatments, although it is not clear what direction the difference will take on aggregate, if any. Observing a lead bid that is close to one's own, in SLB, could encourage a bidder, that might otherwise exit, to stay in the auction. The opposite can also happen. A bidder that would not have exited the auction under RO might be pushed to do so if they observe a lead bid that is too far away. We fall back on the predictions of the SAWOC model reported in Table 2 to formulate the following hypothesis:

Hypothesis 1: *Relative frequency of exit is expected to be higher in SLB compared to RO.*

Regarding any differences across bid compression conditions, in RO we do not expect to find these to have any correlation with exit behavior. For SLB we already noted when studying the simple two bidder model (see Remark 3) that we expect more exits when bid compression is low. This is also reflected in the SAWOC model's predictions.

Hypothesis 2: *Relative frequency of exit remains stable across bid compression conditions in RO but is higher under low bid compression in SLB. With high bid compression the relative frequency of exit is the same in RO and SLB.*

Our intuition regarding auction performance in each treatment and depending on the level of bid compression relies mainly on the SAWOC model's predictions. We expect the relative frequency of efficient outcomes to be lower in RO compared to SLB. The additional information provided in SLB should allow all bidders form more accurate beliefs, which would help bidders with the lowest cost to avoid 'mistake exits' and eventually win the auction.

Hypothesis 3: *The relative frequency of efficient outcomes is higher in SLB compared to RO.*

In RO 'mistake exits' should be more common conditional on bid compression being high: when this is the case the rank of initial offers are less correlated with the rank of bidders' cost. This makes it more likely for the bidder with the lowest cost to be ranked low initially and decide to exit. In SLB, the opposite should be true. Observing a lead bid that is close should encourage all bidders to stay, including the one that has the lowest cost. Therefore, efficient outcomes should be more frequent when bid compression is high.

Hypothesis 4: *The relative frequency of efficient outcomes in RO is higher when bid compression is low. The opposite is true in SLB.*

The SAWOC model predicts that winner's profits will be substantially higher in SLB compared to RO. Notice that what drives this difference are not jump-bids or any 'sealed-bid effect', as there are no dynamics in the SAWOC model. Instead, it can be attributed entirely to differences in exit behavior. What frequently happens in the model in SLB is that by observing the lead bid, the bidder with the 2nd lowest cost correctly updates his belief about the chance of winning the auction downwards and chooses to exit. Thus, the bidder with the lowest cost wins the auction, but at a price that is substantially higher, as the 2nd lowest cost bidder (and likely the others) choose to exit.

Hypothesis 5: *The average Winner's profit is higher in SLB compared to RO.*

Even though in the SAWOC model the average exit rate is similar in RO and SLB when bid compression is high, there are differences in the distribution of exits. Suppose that the bidder with the lowest cost does not exit. In RO it is always the two bidders with the highest initial offer that exit. This is not very likely to affect the price paid to the winner, as it must be the case that the 2nd lowest cost bidder is one of them. Given the correlation of initial offers with the bidders' cost, this happens rarely, although somewhat more frequently when bid compression is high. Overall in RO the average winner's profit is about 1/3 higher than the one expected without exits (=166.66). In SLB exit behavior tends to be more extreme: either no one exits or all but the 1st ranked bidder exit, i.e. a "rush". Intermediate cases can also happen but less often. When bid compression is high, the two cases happen at approximately equal frequencies, therefore the relative frequency of exits overall is similar to the one in RO. Still, whenever there are such "rushes", the bidder with the lowest cost wins and receives a very high price.

Hypothesis 6: *In RO average winner's profit is only slightly higher when bid compression is high. In SLB winner's profit is substantially higher when bid compression is low. When bid compression is low, average winner's profit is higher in SLB compared to RO.*

It should be emphasized again that in formulating these hypotheses we rely mainly on the static models we discussed and ignore any role of auction dynamics. The expected outcomes we describe depend mainly on bidders' exit choices, which in turn we assume that they depend solely on what bidders learn in the beginning of the auction. As we show in what follows, while the intuition build from these static models can help to partly understand what we observe in the experiment, we find significant departures from these predictions that we can attribute to a large extend to the way the bidders in the experiment learn from the feedback they receive during the auction.

We are now ready to present our experimental results, which we do in the following section.

4. Experimental Results

4.1 Aggregate outcomes: Exit behavior

We first examine exit behavior. The first two rows of **Table 3** report the relative frequency of exit in the first three second of every auction. In this initial short period behavior is expected to be mainly driven by the pre-auction information available to bidders. This makes these numbers more comparable to the ones from the SAWOC model reported in **Table 2**. We find that in both treatments bidders exit the auction much less frequently than predicted in this initial period. There are no significant differences across treatments and bid compression plays no role.

Nevertheless, the dynamics of exit behavior start to diverge across treatments as the auction proceeds. In **Figure 1** we graph the estimated survival function (the probability that a bidder remains active in the auction at a given time if they have not exited the auction up to that point) for bidders in each treatment. As can be seen, overall there is more bidder exit in SLB compared to RO. Also, bidders in SLB start exiting sooner, although the rate of exit decreases towards the final third of the auction.

		All data	**Low bid compr.**	**High bid compr.**
Exit *(first 3sec)*	*RO*	.080	.080	.080
	SLB	.086	.087	.085
Efficient Outcomes	*RO*	.693	.685	**.704**
	SLB	.678	.753	**.585**
Winner's Payoff	*RO*	**147.5**	**129.1**	170.6
	SLB	**184.1**	**207.6**	154.7

Table 3: *Average outcomes in the experiment.* The first column is based on all the data. The second and third column only use data from auctions where bid compression was low and high respectively. The first two rows show the relative frequency of exit. The next two show the average time of exit measured in seconds from the start of the auction. The next two show the relative frequency of efficient outcomes. The last two report the average winner's profit. Numbers in bold indicate statistically significant differences across treatments for the corresponding average (*p-val<.05*) in a Wilcoxon matched-pairs signed-rank test where we use group averages as the independent observations.

Result 1: *We find support for Hypothesis 1. The relative frequency of exit in higher in SLB compared to RO, although the difference can only be observed after a few initial seconds of the auction.*

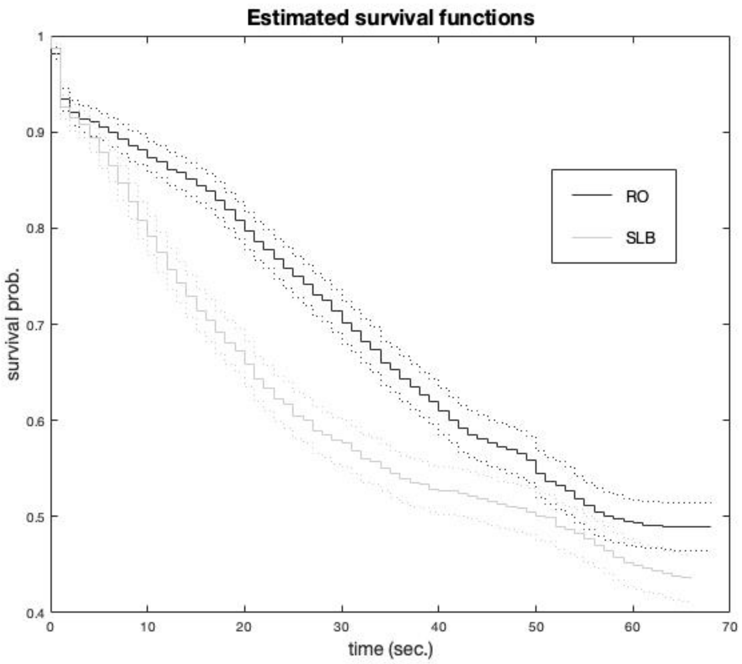

Figure 1: *Estimated survival functions in each treatment.* The two solid lines represent the estimated survival function for bidders in RO (black) and SLB (gray). The dotted lines next to each line represent 95% confidence intervals for the survival probability at each point in time.

As expected, when conditioning on initial bid compression in RO we find no differences in exit behavior across the two conditions of high vs. low bid compression. Interestingly, and contrary to our expectation, we also do not find any differences when doing the same in SLB (see **Figure 4** in the Appendix)

Result 2: *We find no diff*erences *in exit behavior when conditioning on initial bid compression within each treatment. Hypothesis 2 is therefore only supported for RO, but not for SLB.*

Having said that, we shall see in the next subsection that the relative frequency of exit for SLB reported here hides some important heterogeneity in terms of the characteristics of the bidders exiting at any given point. More generally, the dynamics of exit behavior will account for the treatment effects on auction performance that we report below, and the way they differ from our predictions.

4.2 Aggregate outcomes: Efficiency and Winner's profits

The 3rd and 4th rows of **Table 3** report the average frequency of efficient outcomes in each treatment. When comparing the two treatments overall (first column), we do not find significant differences in the frequency of efficient outcomes. In fact, contrary to our prediction, efficient outcomes are slightly more frequent in RO compared to SLB. In general, allowing bidders to exit opens up the possibility of the bidder with the lowest cost to mistakenly exit the auction. It was therefore expected that full efficiency will not be achieved (see **Table 2**). But despite staying in the auction longer than anticipated, bidders in the experiment are not able to achieve even the level of efficiency predicted by the SAWOC model, in neither of the two treatments.

Result 3: *We do not find any support for Hypothesis 3. There are no differences on the relative frequency of efficient outcomes across treatments.*

According to Hypothesis 4, efficiency in RO should be lower when bid compression is high, as it is then that the bidder with the lowest cost might be initially ranked low, negatively update their belief about their chances of winning, and thus exit. This does not turn out to be the case in the experiment.

In fact we observe no significant differences in the frequency of efficient outcomes in RO when conditioning on bid compression. The exact reverse is true for SLB: we expected higher efficiency with high bid compression but find that it is that case that we obtain the fewer efficient outcomes. One thing to notice is that efficient outcomes are significantly lower when bid compression is low in SLB compared to the same case in RO. It turns out that these inefficient outcomes in SLB mostly concern cases where the winner's cost is not very far from that of the bidder with the lowest cost.[22]

Result 4: *Contrary to Hypothesis 4, there is no negative correlation between bid compression and efficiency in RO, while the opposite is true in SLB.*

On the other hand, the two treatments do seem to produce significantly different results with regards to the bidders' payoffs. As evidenced by the measurement of the winner's profit, bidders do achieve on average a higher profit in SLB, compared to RO. This comes at the expense of the buyer, that pays on average a higher price in SLB. Hence, showing the lead bid works in favour of the bidders and against the buyer.

Result 5: *We find support for Hypothesis 5. Bidder's profit is higher in SLB compared to RO.*

The picture becomes more nuanced when we breakdown the samples according to initial bid compression. We find that the treatment effect of SLB on winner's profits seems to derive almost exclusively from the cases with a low bid compression. It is in these cases where the winning bidder achieves on average a substantially higher payoff in SLB compared to RO. Within each treatment, as predicted, high bid compression is associated with higher bidder profits in RO and lower bidder profits in SLB, although the differences are not significant in RO. Notice, that while the SAWOC model correctly predicts the comparative statics across treatments and conditional on bid compression, it substantially exaggerates the expected bidder's profits in the SLB treatment. These predictions rely on bidders in many instances in SLB exiting 'en masse'. While the exit rate in SLB is indeed higher, especially in the earlier stages of the auction, bidders still stay long enough for prices to be driven down substantially.

Result 6: *We find partial support for Hypothesis 6. When bid compression is high, winner's profit is higer in RO, but the difference is not statistically significant. When bid compression is low, winner's profit is significantly higher in SLB.*

A high winner's profit reflects an increased spending on the part of the buyer. In fact, the price paid by the buyer could be higher than expected even in cases where the bidder's profit is low.[23] While we do not report any measure of the buyer's spending here, this turns out to be highly correlated with the winner's profit in our experiment.

Overall, while some of our hypotheses are confirmed by the data, we do find substantial deviations from what is suggested by our theoretical benchmarks. It is worth mentioning that the practitioners in the team of authors had anticipated these results to a large extend, based on their experience in applying auctions in real procurement scenarios. Of course, the theoretical benchmarks we construct completely ignore the dynamic nature of the auction, how bidders may update their beliefs when observing what is happening, and how that can affect their behavior and the auction's performance. In what follows, we take a closer look at what happens during the auctions in the experiment to shed some light on these topics.

---

[22] This can be seen when we measure the excess cost of an allocation, i.e. the difference between the winner's cost and the lowest cost. It turns out that the differences between treatments, overall or conditioning on bid compression, are not statistically significant.

[23] This would happen, for example, if the bidder with the lowest cost and 2nd lowest costs exit, the bidder with the 3rd and 4th lowest cost do not exit, their costs are very close together, and the former wins with an offer very close to the cost of the latter. The price the buyer pays is much higher than expected in the equilibrium without exit, but the winner's profit is low.

4.3 Learning and competition

While we do not directly observe or elicit bidders' beliefs during the auction, their exit behaviour may reveal something about how they learn. Under complete information, any bidder except for the one with the lowest cost would be better off exiting the auction. Thus, if bidders are learning during the auction, we should expect exits to become more correlated with a bidder's *cost rank* (how their cost is ranked compared to that of the others) as the auction progresses. At the beginning of the auction they only have a noisy signal about their cost rank in the form of their exogenously given initial offer's ranking (in RO) or their initial offer's ranking and the leading bid (in SLB).

The regressions shown in **Table 4** demonstrate that in the RO treatment, at the start of the auction (first column) the bidders' exit decision was at least partly explained by their initial ranking, but not by their cost ranking. Once the auction starts (next three columns), bidders seem to gradually learn more about their true cost ranking, resulting in their exit decision being increasingly correlated with their cost rank. The initial ranking does not seem to explain exit decisions after the initial phase of the auction. A similar picture emerges in SLB, only that in this treatment learning appears to be much faster. Already for exit decisions in the first 20 seconds of the auction, the coefficient for cost rank reaches the same level it reaches in the last 20 seconds in RO.

dependent var: *exit* in **RO**	***< 3***	***3-20***	***20-40***	***40-60***
initial rank	.155 (.019)	.041 (.160)	.042 (.022)	.028 (.606)
cost rank	.036 (.145)	.154 (.038)	.285 (.003)	.340 (.001)
constant	-2.018 (.001)	-1.773 (.000)	-1.737 (.000)	-1.923 (.001)
# Obs.	*1600*	*1472*	*1291*	*987*
LogLik	*-433.4*	*-532.7*	*-650.2*	*-445.5*

dependent var: *exit* in **SLB**	***< 3***	***3-20***	***20-40***	***40-60***
initial rank	.071 (.103)	.029 (.449)	-.017 (.381)	.055 (.486)
cost rank	.090 (.063)	.402 (.001)	.397 (.003)	.432 (.000)
constant	-1.875 (.000)	-2.011 (.000)	-1.909 (.000)	-2.434 (.000)
# Obs.	*1600*	*1463*	*1076*	*846*
LogLik	*-459.1*	*-730.2*	*-489.0*	*-293.6*

Table 4: *Learning during the auction.* Probit regressions showing how the choice to exit depends on a bidder's initial offer rank and their cost rank at different stages during an auction. The different columns show the estimated coefficients using data from different time periods/phases of the auction: the first 3 seconds, seconds 3 to 20, 20 to 40 and 40 to 60. The parentheses under the coefficients show the p-values obtained using the wild cluster bootstrap, with robust errors clustered at the group level (see Cameron et al., 2008 ; MacKinnon and Webb, 2018 ; Roodman et al., 2019).

The regression results shown here indicate that learning occurs to some degree in both treatments. Showing the lead bid accelerates this process. We find further evidence to support this conclusion when graphing the survival functions for bidders in each treatment conditional on their cost rank. **Error! Reference source not found.** graphs the survival functions of bidders conditional on their cost rank. The correlation between exits and a bidder's cost rank is clearly visible in these graphs. By the end of the auction, bidders' exit rates are ordered the same way as their cost rank. While this ordering is pretty clear in SLB, in RO we observe that the differences in survival probabilities across bidders conditional on their cost rank are less pronounced. For instance, bidders ranked 3rd and 4th have a very similar survival rate.

Together with the regression results, these graphs seem to indicate that bidders are, to some degree, able to figure out their true chances of winning during the auction and decide upon whether or not to exit accordingly. The richer feedback they receive in SLB, allows them to learn this *faster*. It also seems to be the case that in SLB they learn about their chances more *accurately*. To see this, we can look at the bidders with the lowest cost that should stay in the auction and not exit. As can be seen, they exit at a slightly higher rate in RO compared to SLB. Similarly, from the remaining bidders, except from the ones with the second lowest cost that have similar survival probabilities in both treatments, all other bidders exit at higher rates in SLB. This can be viewed as a manifestation of the discouragement effect induced by observing the lead bid. In SLB, bidders with the highest cost quickly realize they have very low chances of winning the auction and choose to exit.

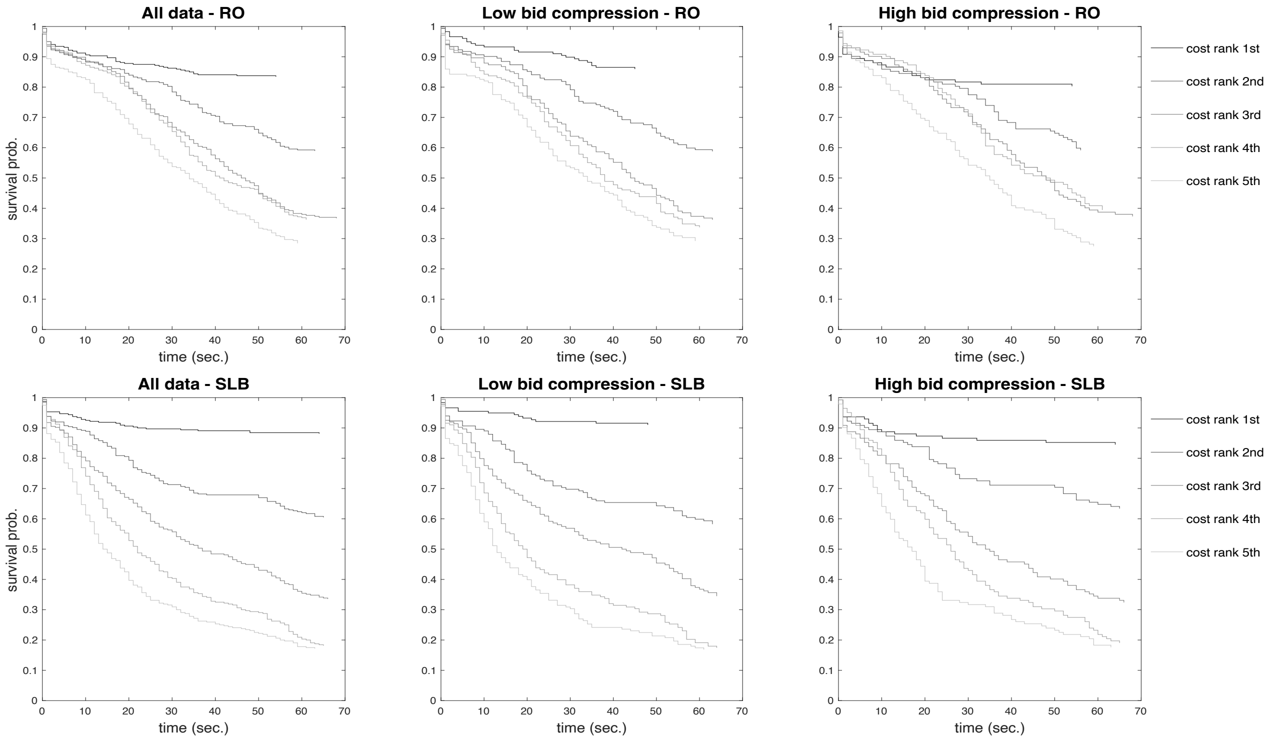

Figure 2: *Estimated survival functions conditional on cost rank.* The lines in each panel represent the survival function of bidders conditional on the rank of their cost, from lowest cost (1st - black) to the highest cost (5th – light gray)). Top panels corrspond to RO treatments and lower panels to SLB. Panels on the right use all the data in each treatment, panels in the middle use data from auctions with low bid compression, and panels on the right use data from auctions with high bid compression.

The graphs also elucidate the role of bid compression. By comparing the panels corresponding to low and high bid compression in each treatment, we observe that high bid compression is associated with a smaller correlation between cost rank and survival probabilities in the first seconds of the auction. Continuing our interpretation of exit rates reflecting bidders beliefs about their chances of winning, it appears to be the case that when bid compression is high, bidders beliefs are less accurate, apparently influenced by the initial offer's rank information. Recall that when bid compression is high, the initial offer's rank is less correlated with a bidder's cost rank (see **Table 5** in the Appendix).

One effect this has is that bidders with the lowest cost tend to exit more (compared to when bid compression is low), with this being more pronounced in RO. A second, and we believe more important effect, is that in SLB with low bid compression, bidders with the second lowest cost exit at a substantially higher rate.[24] As a result, the level of *effective competition* in the auction drops substantially. It is the presence of the bidder with the second highest cost, along with the one with the lowest cost, that can give rise to competitive bidding that can drive prices down. If this bidder exits early, and unless the next bidder's cost is very close, there will not be sufficient competition during the auction for the bidder with the lowest cost to push the price down. This again a manifestation of the discouragement effect in SLB. Notably, for the bidder with the 2nd lowest cost discouragement is significant only in the case of low bid compression.

To see how prices evolve during the auction we measure the difference of the leading bid from the initial lower bid. For comparability across the auctions, we normalize this by the difference of the initial leading bid to the lowest cost for each auction. In **Figure 3** we show the average of this normalized price and it evolves during the auction. The final price levels reflect the differences that we have already observed in **Table 3** regarding winner's profits.

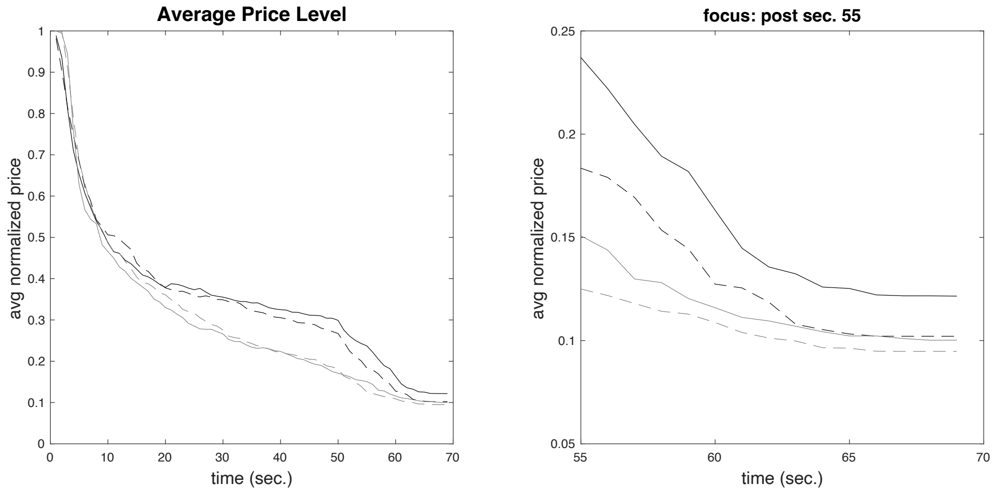

Figure 3: *Price Level.* The panel on the left shows the evolution of the average difference between the leading bid and the lowest cost, normalized by the difference between the initial leading bid and the lowest cost, in each treatment and conditional on initial bid compression. The panel on the right zooms in to the final 15 seconds. Light gray lines correspond to RO and dark gray to SLB. Solid lines correspond to low nid compression and dashed lines correspond to high bid compression.

Overall, we can see how the lower exit rates in RO translate into higher effective competition, which in turn drives prices down much faster than in SLB. After the second 50 in SLB, we observe a new phase of the price dropping. The panel on the right allows us to better distinguish what happens in this phase. Prices in SLB eventually drop to levels that are comparable to those in RO. Nevertheless, when the initial bid compression is low in SLB, the final prices remain at a markedly higher level (~20%) than in the other cases.

Our previous analysis provides an explanation for how these differences in the initial conditions interact with the feedback regime to finally lead to the observed difference in the final price level. In RO, learning is slow and bidders stay in the auction longer. This maintains a sufficiently high level of competition that drives prices down. In SLB, bidders can update their beliefs very fast, and when initial bid compression is low, they do so accurately. This leads to a high exit rate early on, for all bidders except the one with the lowest cost. Conditional to the initial bid compression being low, once the auction reaches its final phase it is very likely for the bidder with the 2nd lowest cost to have exited. When this is the case, there is not enough pressure on the bidder with the lowest cost to reduce the price much further.

[24] In Figure 5 the Appendix show the same data as in Figure 2, but only for the two better ranked bidders, to highlight the differences we mention in the text.

5. Discussion and conclusions

We have focused on the effects of feedback in reverse e-Auctions in the presence of opportunity costs from participating. The option for bidders to exit an auction early opens up the possibility for allocative inefficiencies. This may happen if the bidders that should win the auction get discouraged early and leave early. While we do observe such inefficiencies in our experiment, these are not substantially affected by the type of feedback provided. What is affected are the prices resulting from the auction and the corresponding bidders' profits, although the direction of the effect depends on pre-auction conditions, namely initial bid compression.

A type of discouragement effect seems to be at play. The increased feedback in the Show Lead Bid regime leads to fast learning and a large number of bidders exiting early. But the composition of the exiting bidders is different, depending on the initial bid compression. When bid compression is low, early exit rates are high for all but the most competitive bidders (i.e., the ones with the lowest cost). This leads to an efficient allocation, but the resulting low levels of competition keep prices at a high level, benefiting the winning bidders at the expense of the buyer. When bid compression is high, early exit rates are lower also for bidders that do not have the lowest cost but are still competitive (e.g. the ones with the second lowest cost). While competition in the early stages also drops fast compared to the Rank Only regime, a substantial level of competition persists until the end of the auction, leading to prices comparable to the ones with Rank Only feedback.

Such a mechanism is also compatible with the slight drop in efficiency observed in the SLB treatment with high bid compression: the frequency of efficient outcomes drops, but in terms of excess cost, the loss in efficiency is not substantial. This can be the outcome of a bidding war between the two bidders with the lowest costs, in which sometimes the bidder with the second lowest cost wins. Such "accidents" can be the result of some residual "sniper bidding" behavior in our experiment. Even though the experimental design incorporates a random stopping time to discourage exactly this type of behavior, this solution is not perfect. Typically, procurement auctions use rules extending the auction time if a bid is submitted in the last few minutes of the auction.

For researchers interested in auction design, either theoretically or experimentally, our results highlight how the consideration of an opportunity cost from remaining active in an auction can have a non-obvious effect on its performance. Of course, in auctions for large multi-million contracts any opportunity cost is likely to be negligible compared to the potential profits for a supplier. Still, a very large, and continuously increasing, number of procurement auctions concern more modest contract sizes. In these cases, the opportunity cost introduces elements of all-pay auctions into the environment, affecting bidders' behavior. Further research providing a better understanding of the role of opportunity cost will be welcomed by procurement practitioners that run auctions on a daily basis.

What can these practitioners take away from our research? What feedback regime should be used in procurement auctions where opportunity cost is a concern?

One interpretation of our results is that Rank Only feedback is more robust to such considerations. Bidders stay more engaged, and the resulting competition drives prices down and yields relatively efficient outcomes. The strength of the competitive forces is not affected by pre-auction conditions, namely bid compression. That said, used in practice the Rank Only format may still have some drawbacks. Bidders may dislike the format as it leaves them relatively uninformed, exposes them to the risk of excessive underbidding, and the losers leave the auction without learning much about their competitiveness compared to other market participants. Furthermore, in some instances, such as public procurement in some jurisdictions, showing the lead bid might be a legal requirement.

In light of our results, an alternative approach might be to choose the feedback regime based on the pre-auction conditions in every instance. In fact, this is the advice given to practitioners in relevant handbooks:

"Be aware of the distance between suppliers before the e-auction begins. If initial RFQ bids differ by more than 10 per cent between the lead bidder and the one in second place (especially if it is a winnertakes-all award strategy), a different format should be considered."

(A Practical Guide to e-Auctions for Procurement, Larsen, 2021)

Such an approach requires a good understanding of the market environment and participating bidders, in order to judge what the relevant threshold is in every instance and what other considerations may play a role in determining the design choice. It potentially also exposes the buyer to the "informed principal" problem.[25] Suppliers may use the choice of the design to make inferences about the conditions in the auction. Finally, a frequently changing auction format is not conducive in building trust between suppliers and the buyer. These types of considerations are, of course, only relevant in cases where frequent and repeated participation of the same suppliers is expected. But this is often the case, especially with smaller contract sizes where, as mentioned, opportunity costs are expected to be more relevant.

Another possibility is to develop new designs. One suggestion would be to use a hybrid feedback regime: provide Rank only feedback for an initial period and switch to showing the lead bid after that. Based on what we observe in our experiment, this could potentially maintain competition for the initial phase of the auction driving prices down and keeping bidders engaged. When the lead bid becomes visible offers should already be substantially compressed. Bidders that have stayed behind in their offers might abandon the auction at this point. Still, competitive bidders would continue bidding and drive prices to a desired level. Whether these conjectures have any merit is left for future research.

REFERENCES

Adomavicius, G., Gupta, A., & Sanyal, P. (2012). Effect of information feedback on the outcomes and dynamics of multisourcing multiattribute procurement auctions. Journal of Management Information Systems, 28(4), 199-230.

Anderson, L. R., & Stafford, S. L. (2003). An experimental analysis of rent seeking under varying competitive conditions. *Public Choice*, *115*(1), 199-216.

Ashkenazi-Golan, G., Tsodikovich, Y., & Viossat, Y. (2023). I want to tell you? Maximizing revenue in first-price two-stage auctions. Economic Theory, 76(4), 1329-1362.

Aycinena, D., Bejarano, H., & Rentschler, L. (2018). Informed entry in auctions. International Journal of Game Theory, 47(1), 175-205.

Aycinena, D., & Rentschler, L. (2018). Auctions with endogenous participation and an uncertain number of bidders: experimental evidence. Experimental Economics, 21(4), 924-949.

Babaioff, M., Dobzinski, S., & Oren, S. (2022). Combinatorial auctions with endowment effect. Games and economic behavior, 136, 255-273.

Baik, K. H. (1994). Effort levels in contests with two asymmetric players. *Southern Economic Journal*, 367-378.

[25] See Myerson's (1983) seminal contribution. Specifically for auctions, see for example Jullien & Mariotti (2006) and Skreta (2011).

Baye, M. R., Kovenock, D., & De Vries, C. G. (1993). Rigging the lobbying process: An application of the all-pay auction. *The American Economic Review*, *83*(1), 289.Boone, J., Chen, R., Goeree, J. K., & Polydoro, A. (2009). Risky procurement with an insider bidder. Experimental economics, 12, 417-436.

Boone, J., & Goeree, J. K. (2009). Optimal privatisation using qualifying auctions. The Economic Journal, 119(534), 277-297.

Bruno, G., Esposito, E., Genovese, A., & Passaro, R. (2012). AHP-based approaches for supplier evaluation: Problems and perspectives. *Journal of purchasing and supply management*, *18*(3), 159-172.

Campbell, C. M. (1998). Coordination in auctions with entry. Journal of Economic Theory, 82(2), 425-450.

Cao, X., & Tian, G. (2007). Second-price Auction with Two-Dimensional Private Information on Values and Participation Costs.

Cao, X., & Tian, G. (2010). Equilibria in first price auctions with participation costs. Games and Economic Behavior, 69(2), 258-273.

Chen-Ritzo, C. H., Harrison, T. P., Kwasnica, A. M., & Thomas, D. J. (2005). Better, faster, cheaper: An experimental analysis of a multiattribute reverse auction mechanism with restricted information feedback. Management science, 51(12), 1753-1762.

Deck, C., & Sheremeta, R. M. (2012). Fight or flight? Defending against sequential attacks in the game of siege. *Journal of Conflict Resolution*, *56*(6), 1069-1088.

Duch, M. L., Grossmann, M. R., & Lauer, T. (2020). z-Tree unleashed: A novel client-integrating architecture for conducting z-Tree experiments over the Internet. Journal of Behavioral and Experimental Finance, 28, 100400.

Ehrhart, K. M., Ott, M., & Abele, S. (2015). Auction fever: Rising revenue in second-price auction formats. Games and Economic Behavior, 92, 206-227.

Elmaghraby, W. (2007). Auctions within e-sourcing events. Production and Operations Management, 16(4), 409-422.

Elmaghraby, W. J., Katok, E., & Santamaría, N. (2012). A laboratory investigation of rank feedback in procurement auctions. Manufacturing & Service Operations Management, 14(1), 128-144.

Engelbrecht-Wiggans, R. (1993). Optimal auctions revisited. Games and Economic Behavior, 5(2), 227-239.

Engelbrecht-Wiggans, R., & Katok, E. (2005). Experiments on auction valuation and endogenous entry. Advances in Applied Microeconomics: A Research Annual, 13, 169-193.

Ertaç, S., Hortaçsu, A., & Roberts, J. W. (2011). Entry into auctions: An experimental analysis. international Journal of industrial Organization, 29(2), 168-178.

Fang, D., Noe, T., & Strack, P. (2020). Turning up the heat: The discouraging effect of competition in contests. *Journal of Political Economy*, *128*(5), 1940-1975.

Fonseca, M. A. (2009). An experimental investigation of asymmetric contests. *International Journal of Industrial Organization*, *27*(5), 582-591.

Gradstein, M. (1995). Intensity of competition, entry and entry deterrence in rent seeking contests. *Economics & Politics*, *7*(1), 79-91.

Green, J., & Laffont, J. J. (1984). Participation constraints in the Vickrey auction. Economics Letters, 16(1-2), 31-36.

Hartley, J. L., Lane, M. D., & Hong, Y. (2004). An exploration of the adoption of e-auctions in supply management. *IEEE Transactions on Engineering management*, *51*(2), 153-161.

Haruvy, E., & Katok, E. (2013). Increasing revenue by decreasing information in procurement auctions. Production and Operations Management, 22(1), 19-35.

Hernando-Veciana, Á., & Michelucci, F. (2018). Inefficient rushes in auctions. Theoretical Economics, 13(1), 273-306.

Heyman, J. E., Orhun, Y., & Ariely, D. (2004). Auction fever: the effect of opponents and quasi-endowment on product valuations. Journal of interactive Marketing, 18(4), 7-21.

Hur, D., Hartley, J. L., & Mabert, V. A. (2006). Implementing reverse e-auctions: A learning process. *Business Horizons*, *49*(1), 21-29.

Ivanova-Stenzel, R., & Salmon, T. C. (2004). Bidder preferences among auction institutions. Economic Inquiry, 42(2), 223-236.

Ivanova-Stenzel, R., & Salmon, T. C. (2011). The high/low divide: Self-selection by values in auction choice. Games and Economic Behavior, 73(1), 200-214.

Jullien, B., & Mariotti, T. (2006). Auction and the informed seller problem. *Games and Economic Behavior*, *56*(2), 225-258.

Kagel, J. H., Harstad, R. M., & Levin, D. (1987). Information impact and allocation rules in auctions with affiliated private values: A laboratory study. Econometrica: Journal of the Econometric Society, 1275-1304.

Kagel, J., Pevnitskaya, S., & Ye, L. (2008). Indicative bidding: An experimental analysis. Games and Economic Behavior, 62(2), 697-721.

Klemperer, P. (2002). What really matters in auction design. Journal of economic perspectives, 16(1), 169-189.

Konrad, K. A., & Kovenock, D. (2009). Multi-battle contests. *Games and Economic Behavior*, *66*(1), 256-274.

Larsen, J. G. (2021). *A Practical Guide to E-auctions for Procurement: How to Maximize Impact with E-sourcing and E-negotiation*. Kogan Page Publishers.

Levin, D., & Smith, J. L. (1994). Equilibrium in auctions with entry. The American Economic Review, 585-599.

Li, T., & Zheng, X. (2009). Entry and competition effects in first-price auctions: theory and evidence from procurement auctions. The Review of Economic Studies, 76(4), 1397-1429

Lu, J. (2009). Auction design with opportunity cost. Economic Theory, 38(1), 73-103.

McAfee, R. P., & McMillan, J. (1987). Auctions with entry. Economics Letters, 23(4), 343-347.

Menezes, F. M., & Monteiro, P. K. (2000). Auctions with endogenous participation. Review of Economic Design, 5(1), 71-89.

Milgrom, P. R., & Weber, R. J. (1982). A theory of auctions and competitive bidding. Econometrica, 1089-1122.

Miralles, A. (2008). Intuitive and noncompetitive equilibria in weakly efficient auctions with entry costs. Mathematical Social Sciences, 56(3), 448-455.

Myerson, R. B. (1983). Mechanism design by an informed principal. *Econometrica*, 1767-1797.

Ockenfels, A., & Roth, A. E. (2006). Late and multiple bidding in second price Internet auctions: Theory and evidence concerning different rules for ending an auction. Games and Economic behavior, 55(2), 297-320.

Offerman, T., Romagnoli, G., & Ziegler, A. (2022). Why are open ascending auctions popular? The role of information aggregation and behavioral biases. Quantitative Economics, 13(2), 787-823.

Palfrey, T. R., & Pevnitskaya, S. (2008). Endogenous entry and self-selection in private value auctions: An experimental study. Journal of Economic Behavior & Organization, 66(3-4), 731-747.

Pevnitskaya, S. (2003). Endogenous entry in first-price private value auctions: The self-selection effect. Ohio State Univ., Department of Economics.

Quint, D., & Hendricks, K. (2018). A theory of indicative bidding. American Economic Journal: Microeconomics, 10(2), 118-151.

Reiley, D. H. (2005). Experimental evidence on the endogenous entry of bidders in internet auctions. In Experimental Business Research (pp. 103-121). Springer, Boston, MA.

Skreta, V. (2011). On the informed seller problem: optimal information disclosure. *Review of Economic Design*, *15*(1), 1-36.

Smith, J. L., & Levin, D. (1996). Ranking auctions with risk averse bidders. Journal of Economic Theory, 68(2), 549-561.

Smith, J. L., & Levin, D. (2002). Entry coordination in auctions and social welfare: An experimental investigation. International Journal of Game Theory, 30(3), 321-350.

Stein, W. E. (2002). Asymmetric rent-seeking with more than two contestants. *Public Choice*, *113*(3), 325-336.

Tan, G., & Yilankaya, O. (2006). Equilibria in second price auctions with participation costs. Journal of Economic Theory, 130(1), 205-219.

Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. The Journal of finance, 16(1), 8-37.

Wilson, R. (1977). A bidding model of perfect competition. The Review of Economic Studies, 44(3), 511-518.

Ye, L. (2007). Indicative bidding and a theory of two-stage auctions. Games and Economic Behavior, 58(1), 181-207.

APPENDIX

Supplementary material

Cost rank \ Initial rank	1	2	3	4	5
1	.381	.200	.169	.150	.1
2	.284	.228	.198	.160	.130
3	.171	.196	.222	.215	.130
4	.100	.250	.206	.238	.206
5	.063	.125	.206	.238	.369

Table 5: Cost rank vs. initial rank

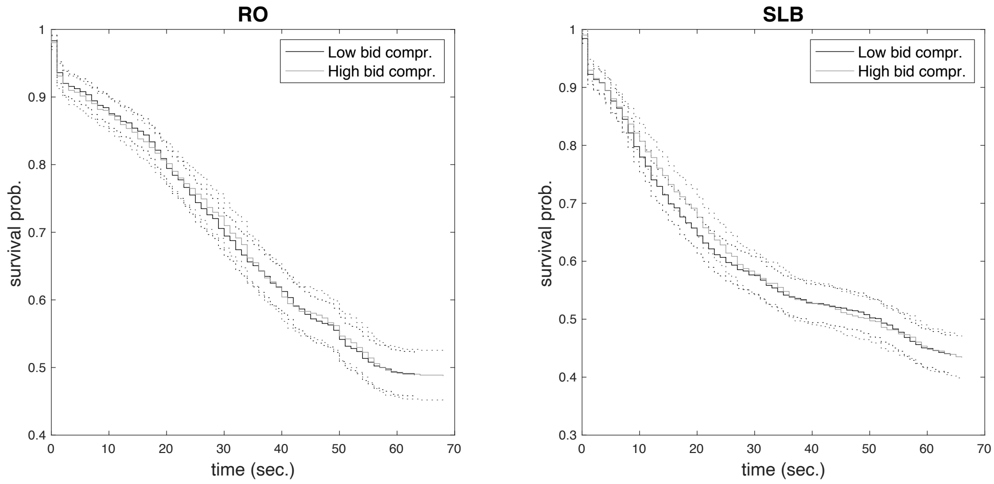

Figure 4: *Estimated survival functions in each treatment conditional on bid compression.* The two solid lines in each graph represent the estimated survival function for bidders in auctions with low (black) and high (gray) bid compression, in the corresponding treatment. The dotted lines next to each line represent 95% confidence intervals for the survival probability at each point in time.

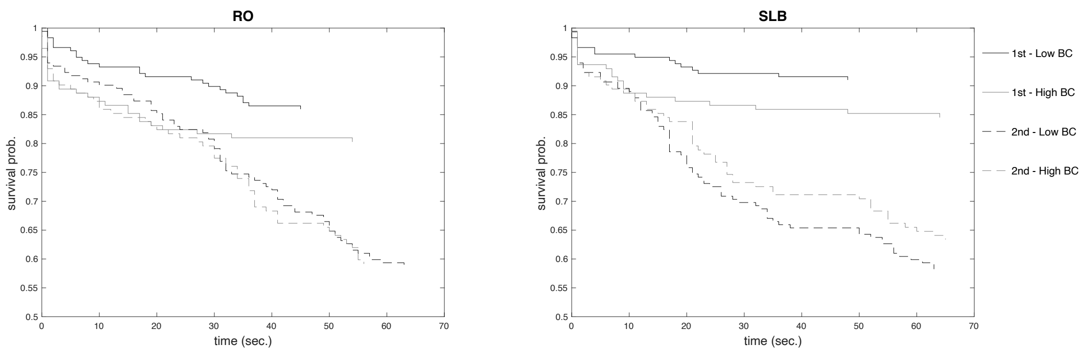

Figure 5: *Estimated survival functions conditional on cost rank for better ranked bidders.* The graphs in the two panels combine top and bottom panels from the middle and right columns in Figure 2, but showing the survival rate only for the bidder with the lowest and the second lowest cost.

Experimental Instructions (translation from Greek)

INSTRUCTIONS

Thank you for participating in this session. Please remain quiet! The entire experiment will be run through the computer and all your decisions will be recorded through it. Please DO NOT talk or make other noises during the experiment. The use of mobile phones and other similar devices is prohibited. Please read the instructions carefully and if you have any question, turn on your microphone to ask.

The task

In this experiment you will have the role of a supplier who competes in each round with others to supply his own product to a buyer. All participants in the experiment will be suppliers. The role of the buyer will be played by a computer that will choose the best offer every time.

There will be a total of 40 rounds in the experiment and the choices and outcomes of each round will be completely independent of the others

Auction

Competition in each round will be through a reverse auction. Each supplier will be able to make offers to the buyer in real-time.

This can be done either by writing the amount of the offer (in the box on the left) or by increasing or decreasing the amount of the previous offer (with the buttons on the right) and pressing the corresponding button 'Submit'. All offers must be integer numbers and refer to points. Each new offer must be less than the supplier's previous valid offer.

The duration of the auction will be 60 + δ seconds, where δ is a number between 1 and 10 that will be randomly selected in each round. No vendor will know exactly when the round ends. A timer will show the time remaining and when it is nearing the end the 'Time is running out!' indicator will flash. Whichever supplier has the lowest bid at the end of the round will supply the product to the buyer. He will pay the supply cost and collect his bid amount. The difference between the two is his profit for that round.

Cost

Each supplier will have a supply cost that they will have to pay if they are selected as a supplier. This cost is determined randomly in each round and independently for each supplier. This will be done as follows: the computer randomly selects an interval of 1000 numbers between 2000 and 8000 (eg {4563,4564,..., 5561, 5562}). From this interval it randomly and with equal probability selects a number that is the cost. For all suppliers the supply cost is selected from the same random interval. So, knowing your supply cost, you know that the other's supply cost can be up to 999 points higher or 999 points lower. Each supplier will know their own costs but not the of others'.

Exit
At any time during a round, each vendor has the option to exit the auction (for that round) and collect some points. The amount he can collect from the exit will start at 30 points at the start of the round and gradually decrease during the round until it reaches zero. If a supplier withdraws, then their current bid is void and cannot be selected as the lowest bid. Exiting a round does not in any way affect participation in previous and subsequent rounds of the experiment.

Ranking
Each supplier will not know the offers of the other suppliers. What he will know will be the ranking of his current offer in relation to the others. At the top of the screen will be listed the ranking of the current offer of each supplier. The best (lowest) bid is 1st, the next highest 2nd, and so on. In this ranking, only the offers from suppliers that have not left in the given round count. The current ranking will be updated in real time.

Initial offer and ranking
For each vendor, at the beginning of each round, an initial bid will be selected in the following manner. The computer will add to each vendor's cost a number between 1500 and 3500 at random and with equal probability. This number will be selected separately and independently for each vendor in each round. The initial bid will also determine the initial ranking of suppliers. By the way initial bids are selected, it follows that if one supplier has a lower initial bid than another, then it is more likely, but not certain, that they also have a lower cost.

Preparation phase
At the beginning of each round there will be a 10" preparation phase. In this phase, each supplier will be able to see on his screen his cost, his initial offer, and his initial ranking. Use this phase to review this information and prepare your strategy for each round.

Earnings
The sum of the points you will collect in each round gives your total earnings from the experiment in points. At the end of the experiment you will be paid 1 euro for every 150 points you have earned. The participation fee of 5 euros will be added to this amount.

Trial rounds
Before the experiment starts we will do a trial run to make sure everyone understands the process. Any "profits" made in the trial round will not count towards the calculation of the total profits from the experiment.