THE UNIVERSITY OF QUEENSLAND
AUSTRALIA
CREATE CHANGE

Discussion Paper Series
UQ School of Economics

Paper Name: A Secret Worth Keeping? Bid Cap Design in Budget-Constrained Procurement Auctions
Date: February 2026
Josephine Auer, Lana Friesen, and Ian A. MacKenzie

Discussion Paper No. 672

A Secret Worth Keeping? Bid Cap Design in Budget-Constrained Procurement Auctions*

Josephine Auer[§,‡], Lana Friesen[§], and Ian A. MacKenzie[§]

[‡]Department of Economics, MIT, Cambridge, MA 02142

[§]School of Economics, University of Queensland, Brisbane, Australia 4072

Abstract

This article investigates the existence of bid caps in budget-constrained procurement auctions. We analyze the design and (non)disclosure of a bid cap and how this impacts aggregate market outcomes and strategic bidding behavior in a budget-constrained environment. We use a laboratory experiment to analyze two potential bid cap designs—a disclosed versus undisclosed bid cap—as well as comparing both to a baseline case without a bid cap. We find adoption of either a disclosed or non-disclosed cap significantly improves cost effectiveness. A non-disclosed cap, however, significantly *increases* the information rent to participants and, consequently, performs relatively worse than a disclosed cap. We consider two common but distinct auction formats (discriminatory 'pay-your-bid' and a uniform price) and show that a discriminatory auction improves cost effectiveness compared to a uniform-price auction when the cap is disclosed. Our findings have important policy implications that demonstrate the benefits of implementing bid caps for improving budgetary cost-effectiveness while highlighting potential tradeoffs between efficiency and worsening information rents.

Keywords: auction, experiment, bid cap, procurement

JEL Classification: C91; C92; D44; H57

*We are grateful for helpful comments provided by Nikhil Agarwal, Glenn Ellison, Drew Fudenberg, Andrew Komo, Claudio Mezzetti, Frank Schilbach, and audiences at the AARES conference (Christchurch), the Australia New Zealand Workshop on Experimental Economics (RMIT), the ESA World Meeting (Lyon), and the Monash Environmental Economics Workshop. We thank the University of Queensland for financial support. The study has approval from the University of Queensland, Business, Economics and Law, Low and Negligible Risk Ethics Sub-Committee (2019/HE001872).

1 Introduction

Procurement auctions are a common process to acquire goods and services, ranging from defense acquisition and electricity generation to conservation and greenhouse gas reduction (Arozamena and Cantillon, 2004; Cramton et al., 2021; Lopomo et al., 2023). A well-known problem in these auctions is that of asymmetric information: a buyer does not know bidders' costs, which may result in significant rent transfer. This problem is exacerbated when the buyer has a finite budget with which to obtain goods and services. For example, in the U.S. Conservation Reserve Program (CRP)—a procurement process to improve multiple conservation areas—Kirwan et al. (2005) estimate that the premium earned by participants constitutes 10-40% of the program's rental payouts. One solution to limit rents is the use of *bid caps*: a prescribed bid level that stipulates a maximum acceptable payment level above which bids will not be considered. This seeks to reduce inflated bids but, equally, may also reduce participation within the scheme. Given the apparent tradeoffs associated with bid caps, this article investigates how the introduction and (non)disclosure of a bid cap affects bidding behavior and cost-effectiveness in procurement auctions.

The implementation of bid caps has evolved over time with the creation of a diverse range of designs associated with the informational requirements of the bid cap. For example, in procurement for the U.S. Department of Defense, an "affordability cap" exists that places a budget constraint on acquisitions, which is unknown to participants (Porter et al., 2015). Equally, the U.S. CRP now discloses bid caps, which previously were not disclosed to participants (Hellerstein, 2017). Within electricity generation, Brazil has a publicly announced price ceiling (Dutra and Menezes, 2005; Tolmasquim et al., 2021) and in Australia's flagship climate policy a bid cap was implemented on the maximum amount payable for carbon dioxide mitigation (denoted as a "benchmark price") but was not disclosed to participants in the first five years but then removed in 2020 so that—officially—no bid cap existed after this date.[1] In the policy sphere, then, a plethora of bid-cap designs exists, but there is less attention paid to the optimal design and implementation of bid caps.

Given the importance of these mechanisms—both in terms of financial burden and potential impact to government cost effectiveness—it is important to understand how the existence and

[1] Although successful bids continued to be subject to sufficient funds being made available. Much discussion has existed over the benchmark price, for example, the Climate Change Authority (2017, p. 71) stated that "*revealing the price gives greater certainty to the market but also may encourage bidding at or just below the benchmark price by scheme participants. The Government decided that the benchmark price would remain confidential but the average price of successful bids would be published after each auction to provide some information on price to future scheme participants and support investment in new projects*".

informational composition of these bid caps influences bidding behavior and the resulting success or failure of the process. In particular, should a bid cap be introduced and, if so, should the bid cap be disclosed to participants or kept secret: is it a secret worth keeping?

A natural question is why a laboratory experiment is informative in this setting. Although many procurement programs generate rich administrative data, real-world variation cannot credibly identify the causal effect of cap disclosure versus secrecy. Policy adoption is endogenous—governments choose whether to impose or reveal a cap in response to expected participation, historical bidding patterns, or budgetary pressures—making observational contrasts difficult to interpret. Moreover, auctions differ across programs in ways that confound inference: procurement objectives, cost distributions, project types, and participation rules all vary systematically with cap policy. Furthermore, participation is endogenous and bidder beliefs are unobservable, so field data cannot recover the strategic bidding responses induced by different cap designs. The laboratory setting overcomes these limitations by holding fixed the cost distribution, competitive environment, and budget constraint. This allows us to isolate how bidders adjust behavior when facing a known cap, uncertainty about an undisclosed cap, or no cap at all.[2]

This article focuses on understanding the impact of implementing a bid cap and whether disclosure or non-disclosure of a bid cap to participants impacts bidding behavior and the market outcome. Our main focus is on the commonly used discriminatory auction with 'pay-your-bid' format for multiple homogeneous units in which the buyer has a finite budget. Unlike procurement processes that seek only to identify a supplier for a fixed quantity of a good or service, we study budget-constrained settings—such as the procurement of electricity generation or carbon dioxide mitigation—where the government determines both the level of supply and its associated cost.[3] In this auction type, bidders submit bids and, if successful, they obtain their bid. We focus on three main cases, where (i) no bid cap exists, (ii) a bid cap is disclosed to participants, and (iii) a bid cap exists but remains undisclosed to participants. We derive the equilibrium bidding strategies and experimentally test both the individual- and market-level outcomes. We also extend our experimental analysis to investigate the impact of bid caps in

[2] Although stylized, laboratory evidence remains informative for procurement settings. Experimental auction studies consistently document that core behavioral patterns—such as bid shading and responses to information and institutional rules—map closely to behavior in applied environments (Kagel and Levin, 2016, 2002). This body of work suggests that the bidding adjustments generated by disclosure and secrecy in our setting are likely to be relevant beyond the laboratory, providing a useful complement to observational procurement data.

[3] Within budget-constrained environments, significant government funds have been invested. For example, additional applications include Australia's Emissions Reduction Fund (ERF) with a A$2.5bn fund and carbon capture and storage (removal) (including the €3bn Bioenergy with Carbon Capture and Storage (BECCS) Fund in Sweden and the Dkr. 166m Negative Emissions via Carbon Capture and Storage (NECCS) Fund in Denmark).

a uniform-price procurement auction, where all successful participants receive the same price irrespective of their chosen bid.

Our main aim in this article is to experimentally compare auction performance when implementing different types of bid caps. Our primary measure of performance is Budgetary Cost Effectiveness (BCE), which measures the number of dollars spent per units purchased, a measure that procurers will want to minimize. For additional insight, we decompose BCE into two opposing effects: efficiency and rent (Schilizzi, 2017). Efficiency reflects the extent to which the outcomes have been achieved in a least-cost manner thus improving BCE. Rent reflects the magnitude of informational rents captured by bidders—the amount spent above the cost—which worsens BCE. This allows us to examine the source of improvements in BCE and to identify any potential tradeoffs—for example, an auction format might improve one component at the expense of the other.

We find, based on the aggregate market outcomes in our experiment, that implementing a bid cap in a discriminatory procurement auction significantly improves BCE regardless of whether the cap is disclosed or not. This improvement results from significant improvements in efficiency that offset increases in information rents received by bidders under both types of caps. Somewhat counterintuitively, the increase in rents is only significant when the cap is undisclosed. Although one might expect secrecy to discipline bids and thereby reduce rents, uncertainty about the cap's level instead distorts bidding, leading to higher overall information rents. This provides evidence that either type of cap works as intended—the net cost effectiveness is improved with only a relatively small rent increase—but keeping the cap secret may not be justified; the cap should be implemented *and* disclosed.

Turning to the individual level, our analysis of experimental bidding behavior provides supporting evidence consistent with these market outcomes. In particular, we find that implementing a bid cap of either kind significantly reduces average bids across the entire cost range, consistent with the improvement in BCE described above. We also find that for moderate cost values, disclosed bid caps significantly reduce bids compared to undisclosed bid caps, consistent with our findings regarding differences in rents.

Extending our experimental analysis to incorporate a different auction format—a uniform-price auction—we find similar evidence that bid caps significantly increase cost effectiveness but—in contrast to a discriminatory auction—we do not find significant individual impacts from either the efficiency or the information rent effect. Comparing the two auctions formats, we find that in the presence of a bid cap, using a discriminatory auction format improves cost effectiveness compared to a uniform-price auction, although this difference is significant only when the bid cap is disclosed. Interestingly, while in the absence of a bid cap we find no significant difference in budgetary cost-effectiveness between the two auction formats, this masks an important tradeoff: while rents are significantly reduced in the discriminatory auction, efficiency is worse compared to the uniform-price auction.

Our overarching contribution is to provide the theoretical and experimental underpinnings of bid caps in contemporary procurement auctions with budget constraints, which helps shed light on their optimal implementation. We show that bid caps—regardless of type—improve cost effectiveness but, equally, there exist clear idiosyncratic impacts on bidding behavior with respect to distinct bid cap design and auction format. In particular, keeping the cap undisclosed appears less justifiable than full disclosure. Further, the auction format—either a discriminatory or uniform-price auction—has significant impacts, with a discriminatory auction performing (weakly) better than a uniform-price auction in terms of budget-cost effectiveness for the full range of bid cap designs. Together, these results motivate the theoretical and experimental analysis that follows, which clarifies when bid caps improve procurement outcomes and how their design shapes behavior in budget-constrained auctions.

1.1 Related Literature

Our findings contribute to the extensive literature on procurement auction design. Procurement auctions have been extensively investigated with a focus on, for example, the impact of investment decisions (Arozamena and Cantillon, 2004), firm entry (Li and Zheng, 2009), bid preferencing (Marion, 2007; Nakabayashi, 2013; Hyytinen et al., 2018; Arozamena et al., 2023), bidding behavior (Chassang et al., 2022), collusion (Chassang and Ortner, 2019; Kawai and Nakabayashi, 2022), information provision (De Silva et al., 2009), quality concerns (Lopomo et al., 2023), and finite (but unknown) budgets (Rottner, 2025). Yet in the study of general procurement auctions, the use of bid caps in a budget-constrained environment—now a central feature of many applied procurement programs—has received surprisingly little systematic attention.[4] This article addresses that gap by examining how alternative bid cap designs shape efficiency and information rents.

Only a limited number of studies investigate bid caps by varying the size of the bid cap in the context of a known procurement target (e.g., a target number of fixed units to be purchased)

[4]Beyond the procurement-specific literature, recent mechanism-design results highlight the broader appeal of pay-as-bid formats. Pycia and Woodward (2025) show that such mechanisms can be optimal in wide classes of allocation environments. Their theoretical insights complement our analysis by providing a foundation for understanding how institutional features common in procurement shape outcomes in pay-as-bid auctions.

rather than a budget-constrained environment, which is the focus within this article.[5] Hellerstein and Higgins (2010) examine different levels of group-specific bid caps in the context of a target-constrained auction with quality adjustments. Cramton et al. (2021) investigate the level of the bid cap, also in the context of a target auction. They both find that setting the bid cap level is a 'balancing act' between reducing participation and restricting rents. These previous studies, therefore, focus on the impact of the *level* of the bid cap with a procurement target. Our work, in contrast, contributes to the literature by comparing the disclosure and non-disclosure of bid caps relative to no bid cap in a variety of budget-constrained auction format settings and are able to provide guidance as to the desirability of implementing and designing a bid cap for procurement.

Our work highlights the appropriateness of using bid caps to improve budget cost effectiveness in current programs that employ them—such as environmental, defense, and energy procurement—and provides guidance for their implementation in other contexts. For program designers, we provide evidence that the implementation of a bid cap is justifiable to improve cost effectiveness and non-disclosure of a bid cap is not a necessary condition for this to occur: allowing participants to know the bid cap works equally (if not) better at budget cost effectiveness.

The article is organized as follows. In Section 2 the conceptual theoretical framework is derived. In Section 3 the experimental framework is outlined. Section 4 details the experimental results and Section 5 extends the main framework to consider a uniform-price procurement auction. Section 5.2 compares the outcomes across both auction formats. Section 6 provides some discussion and concluding remarks. Appendices A and B provide the theoretical framework with formal proofs and Appendix C provides additional results and findings. Appendix D provides a sample of the experimental instructions.

2 Basic Model

Consider a procurement auction in which a single buyer—the auctioneer—has a budget of $\bar{M} > 0$.[6] The buyer wishes to purchase as many units of some homogeneous good as possible

[5]Secret reserve prices have been studied in standard auctions yielding ambiguous results. For example, Riley and Samuelson (1981), Elyakime et al. (1994), and Rosenkranz and Schmitz (2007) show that bidder revenue is maximized by announcing a public reserve price. On the other hand, in the context of a common-value second-price auction for a single unit, Vincent (1995) provides an example of how a secret reserve price can, in fact, increase ex-ante expected profits relative to a case where the reserve price is announced. Note while procurement auctions have many similarities to standard auctions, such as pollution auctions, analysis requires a unique approach due to the procurement nature of the activity (such as the addition of bid caps and fixed budgets). For an overview on pollution auctions see Betz et al. (2017), MacKenzie (2022), and Friesen et al. (2022).

[6]The theoretical analysis can be extended to the case with no binding budget constraint, letting $M \to \infty$.

with this budget. Suppose a risk-neutral bidder $i \in \{1, 2, \ldots, n\}$ can sell the good by expending a marginal cost c_i for all units up to some common capacity $\bar{\lambda}$. For ease of exposition, we normalize $\bar{\lambda} = 1$ and assume each bidder submits a single per-unit bid that applies to all units.[7] We suppose that each bidder can submit only whole units (i.e., no partial sales), which aligns with common programs, including the U.S Conservation Reserve Program (CRP) and Australia's Emissions Reduction Fund. The marginal cost of bidder i is privately known and independently drawn from a common cumulative distribution $F(c_i)$ with probability density function $f(c_i)$ on the support $[\underline{c}, \bar{c}]$.

Suppose that a per-unit bid cap $K \geq 0$ exists that limits the maximum value of accepted bids from bidders for each unit. Our main focus of this article is to analyze the impact of bid caps. We focus on cases where (i) no bid cap exists (ii) a bid cap exists and is disclosed to bidders and (iii) a bid cap exists but is undisclosed to bidders. When the bid cap is undisclosed, bidders know only that it is a random variable K drawn from a commonly known distribution $H(K)$ with continuous density $h(K)$ on the support $[0, \bar{K}]$ for $\bar{K} > 0$.

We focus on the most common auction format for procurement auctions that is experienced in policy settings and consider the discriminatory 'pay-your-bid' auction. In the discriminatory auction, each successful bidder receives payment for each unit equal to the individual offer they submitted in the sealed-bidding process. To determine the successful bidders, the auctioneer first calculates the total value of offers at *each* bid price, using the list of bids sorted in ascending price order. The buyer then purchases units in this order until one of the following two conditions occurs: (i) the auctioneer's fixed budget, $\bar{M}$, is exhausted, or (ii) the relevant bid cap, K, is breached. The model is defined by ex ante *symmetric* bidders, where we assume each bidder $i \in \{1, \ldots, n\}$ follows the symmetric bidding strategy, $\beta : [\underline{c}, \bar{c}] \longrightarrow \mathbb{R}_+$. In Section 5 we extend the analysis to include a uniform-price auction and compare outcomes with respect to the discriminatory auction.

Let b_i denote the bid submitted by bidder i, which determines the minimum price bidder i is willing to accept for the number of units offered, $q_i \in \{0, 1\}$. The aggregate supply is therefore the sum of individual bids and can be represented by a following step function $S(p)$, where p is the price per unit:

$$S(p) = \sum_{i=1}^{n} q_i|_{b_i \leq p}. \tag{1}$$

As discussed above, the final price received by successful bidders in the discriminatory auction,

[7] Given bidders have common capacities and have a single bid for all units, the discriminatory auction can be generalized to $\lambda > 1$ without qualitative changes in bidding strategies. In Appendix B, we show that it is a weakly dominant strategy to submit an offer for all λ.

p_i^*, is determined by the individual bid offers.

To aid in the theoretical and experimental analysis we define the following key measures of auction performance that will be used throughout (Schilizzi and Latacz-Lohmann, 2007; Schilizzi, 2017).

Definition 2.1. For the procurement auction:

- *Efficiency (Eff)* denotes the extent to which a level of supply has been achieved in a least-cost manner, such that
$$Eff = \frac{TAC}{Q},$$
where TAC denotes the aggregate total actual cost to firms and Q is the aggregate supply.

- *Rent* denotes the magnitude of the information rent captured by bidders, such that
$$Rent = \frac{TC}{TAC},$$
where TC is the total cost of procurement to the government.

- *Budgetary Cost Effectiveness (BCE)* is the total cost to the government relative to the level of supply, which can be decomposed into both Eff and $Rent$ effects such that
$$BCE = \frac{TAC}{Q} \cdot \frac{TC}{TAC} = \frac{TC}{Q}.$$

Given these definitions, this article aims to understand how the introduction and alternative informational structures of a bid cap influences efficiency, rent acquisition, and overall budget cost effectiveness of the procurement process. We start by considering a scenario in which no bid cap exists. All proofs are contained in the Appendix.

2.1 No bid cap

Suppose no bid cap exists such that the buyer accepts offers until the budget $\bar{M}$ is exhausted. Define $\bar{m} = \frac{\bar{M}}{n}$ as the balanced and equal payment to bidders if the buyer's budget was equally split between bidders. Given this budget, we outline an individual bidder's equilibrium bidding strategy when no bid cap exists.

Proposition 1. *For a finite budget $\bar{M}$ with a discriminatory auction and no bid cap, bidder i's equilib-*

rium bid is the solution to the differential equation,

$$\beta'(c_i) = (\beta(c_i) - c_i)\frac{\sum_{k=1}^{n-1} f(c_i)G_k(c_i)\left\{\frac{(n-1)F(c_i)-k}{F(c_i)(1-F(c_i))}\right\}}{\sum_{k=0}^{n-1} G_k(c_i)}, \tag{2}$$

where,

$$G_k(c_i) := Pr\left(\sum_{j=1}^{k+1} \beta(c_{(j)}) \leq \bar{M}\right) F(c_i)^k(1-F(c_i))^{n-k-1},$$

and $c_{(j)}$ is the j^{th} order statistics of $c_1, \ldots, c_n$.

Proposition 1 provides an implicit equilibrium bidding strategy for each bidder. In this baseline case, bids exceed costs because bidders face a tradeoff between raising their bid (to earn more conditional on winning) and keeping it low enough to be allocated before the budget binds. The equilibrium markup arises directly from this probability-of-allocation effect.

2.2 Disclosed bid cap

Let us now consider the case where a bid cap exists. We start by considering the equilibrium bidding strategy if the bid cap is disclosed to bidders prior to the auction, denoted with superscript "*DDisc*". Intuitively, as one would expect, if there is a known (binding) bid cap and the buyer's budget is sufficiently large to accommodate all bids up to the cap, then each bidder with lower private costs will submit a bid equal to the bid cap. We refer to such a case as having a *non-binding* budget (i.e., $\bar{m} \geq K$). Conversely, the budget is *binding* when the total value of accepted bids would otherwise exceed the available budget $\bar{M}$.

Proposition 2. *Within a discriminatory auction suppose a bid cap $K > 0$ exists that is known to all bidders prior to the auction and the buyer's budget is non-binding. If $c_i \leq K$, it is a weakly dominant strategy to bid $\beta^{DDisc}(c_i) = K$. If $c_i > K$, it is a strictly dominant strategy to bid $\beta^{DDisc}(c_i) > K$.*

With a disclosed bid cap, bidding below or at the bid cap will ensure the bidders win. This, while potentially increasing participation, will also increase the rents transferred from buyer to bidders. That is, we are again in a setting where *Eff* is optimal but *Rent* achieves the worst case upper bound among cases in which all units are sold.

We now turn to the more general case in which the budget, rather than the bid cap, may become the binding constraint. When the budget is the true binding constraint, the cap ceases to affect best responses; hence bids collapse back to the no-cap strategy unless that strategy would exceed K. Put differently, the equilibrium bid for types below the cap is simply the no-cap bid, truncated above at K.

Proposition 3. *In a context in which the budget may be a stricter bound than the disclosed cap, it is a weakly dominant strategy to bid*

$$\beta^{DDisc}(c_i) = \begin{cases} \min\{\beta^{DNo}(c_i), K\} & c_i < K \\ c_i & c_i \geq K, \end{cases}$$

where $\beta^{D-No}(c_i)$ is the optimal bidding strategy in the no-bid case. That is, for $c_i < K$ bidders behave as in the no-cap equilibrium unless their optimal no-cap bid exceeds K, in which case they bid K.

Intuitively, when the budget is a stricter binding constraint than the disclosed bid cap, the bidder's optimal bidding strategy is to bid identical to the no-bid cap case or at the bid cap, whichever is smaller.

2.3 Undisclosed bid cap

We now suppose that the bid cap is unknown to the bidders. In particular, there is common knowledge that the bid cap is distributed according to a distribution function H with a density function h on the support $[0, \bar{K}]$. Proposition 4 details the equilibrium strategy, which is denoted with superscript "*DUnd*".

Proposition 4. *Suppose a buyer's budget is non-binding within a discriminatory auction with an undisclosed bid cap. If $c_i > \bar{K}$, it is a strictly dominant strategy to bid $\beta^{DUnd}(c_i) > \bar{K}$. If $c_i = \bar{K}$, it is a strictly dominant strategy to bid $\beta(c_i) \geq c_i$. If $c_i < \bar{K}$ then the unique payoff maximizing bid is implicitly defined by,*

$$\beta^{DUnd}(c_i) = \min\left\{\frac{1 - H(\beta(c_i))}{h(\beta(c_i))} + c_i, \bar{K}\right\}, \tag{3}$$

provided that $h'(\beta(c_i))(\beta(c_i) - c_i) + 2h(\beta(c_i)) > 0$, where H is the distribution function of the unknown bid cap, K.

Proposition 4 shows that the equilibrium bidding strategy is either the upper bound of the cap or a markup from the firm's costs that depends on the structure of uncertainty in the cap, whichever is smaller.

As in both of the previous settings, the equilibrium bidding strategy is (weakly) increasing in private cost, and hence achieves optimal *Eff*. However, depending on the bid cap distribution, it may be possible for the auctioneer to improve on the *Rent* outcome. If the commonly known distribution, H, is either sufficiently narrow or the upper bound, $\bar{K}$ is sufficiently low, then the auctioneer may achieve greater *Rent* efficiency relative to the no bid cap and disclosed bid cap treatments.

When the budget is binding (rather than the undisclosed cap) there is no analytical solution. Later, in Section 2.5, we numerically approximate the optimal bidding strategy using an unknown cap in which the budget can bind (as well as estimate optimal bidding strategies for both no cap and disclosed cap cases).

2.4 Illustrative example: closed-form bidding for two players

Because the $n = 2$ setting admits a closed-form equilibrium bidding function, we include this special case as an intuitive illustration of how bid caps shift markups. This example is not intended as a prediction for the experiment (which uses $n = 4$), but instead assists the reader for the numerical $n = 4$ equilibrium derived in Subsection 2.5 for the experimental analysis. Denote with superscript "DNo" the no-bid cap. We have the following closed-form expression for the equilibrium bid with no bid cap.

Corollary 1. *For a finite budget $\bar{M}$ with a discriminatory auction, $n = 2$, and no bid cap, bidder i's equilibrium bid is given by*

$$\beta^{DNo}(c_i) = \begin{cases} \bar{m} & \bar{c} < \bar{m} \\ \frac{1}{1-F(c_i)} \int\limits_{c_i}^{\bar{c}} x f(x) dx & \bar{m} \leq \bar{c} \leq \bar{M} \\ \frac{1}{1-F(c_i)} \left[\bar{M}(1 - F(\bar{M})) + \int\limits_{c_i}^{\bar{M}} x f(x) dx \right] & \bar{M} < \bar{c} \\ c_i & \bar{M} \leq c_i, \end{cases} \tag{4}$$

Overall, we show that the key determinants of a firm's bidding strategy depend on the size of firms' costs relative to the budget and on the distribution of those costs. There are four possible cases. If firms' costs are low relative to the size of the budget (i.e., $\bar{c} < \bar{m}$) so that even the highest possible cost is smaller than the equal split of the budget, then each bidder has an incentive to bid $\bar{m}$. Both bidders win and the budget is fully exhausted. While *Eff* is optimal in this case, *Rent* attains its highest level among cases in which all units are sold. This result suggests an important design consideration for the buyer: when the budget is sufficiently large—one in which it is profitable for all bidders to participate—the buyer effectively dictates the outcome of the auction (and hence the degree of bid inflation) through the choice of the per-bidder budget share $\bar{m}$.

If the buyer is not necessarily intent on purchasing from *all* bidders, it may therefore be more efficient—specifically through a reduction in *Rent*—to set a smaller budget such that $\bar{m} \in [\underline{c}, \bar{c}]$. As the budget tightens within this range, bidders submit their cost plus a strategic mark-up;[8] *Eff* remains optimal, but *Rent* increases relative to a truth-telling benchmark. Corollary 1 formally derives these equilibrium bidding strategies for both budget-sufficient ($\bar{M} \geq \bar{c}$) and budget-constrained ($\bar{M} < \bar{c}$) cases. Finally, when the budget is smaller than a bidder's cost ($\bar{M} \leq c_i$), bidding at cost becomes the best response.

Corollary 2 derives the equilibrium bidding strategy with a uniform distribution of private costs. For uniformly distributed costs, $U[0, \bar{c}]$, when $\bar{M} \geq \bar{c}$, the bidding strategy is as follows.

Corollary 2. *For a discriminatory auction with no bid cap, $n = 2$, and uniformly distributed costs, $U[0, \bar{c}]$, when $\bar{M} \geq \bar{c}$, bidder i's equilibrium bid is given by:*

$$\beta^{DNo}(c_i) = \begin{cases} \bar{m} & \bar{c} < \bar{m} \\ \frac{c_i + \bar{c}}{2} & \bar{c} \geq \bar{m}. \end{cases} \tag{5}$$

If the per-unit budget is sufficiently large to allow every participant to sell a unit then bidders will bid $\bar{m}$. If the per-unit budget is relatively smaller then there is a possibility that if a bidder increases their bid they may lose. So a bidder knows that while increasing a bid will increase payoffs it also decreases the probability of winning. Consequently the bidder's equilibrium strategy is to bid the mid-way point between their private cost and maximum cost: the mark-up is half the distance between their private cost and the maximum cost.

Consistent with the more general case in Proposition 1, the bidding strategy is (weakly) increasing in cost and hence achieves optimal *Eff*. By decreasing the per-unit budget, $\bar{m}$, the auctioneer can improve the *Rent* outcome. Indeed, if the auctioneer could observe private costs, then setting $\bar{m} = \max\{c_1, c_2\}$ would minimize *Rent* among cases in which all units are sold.

For a disclosed cap, the two-player example matches the general logic in Subsection 2.2: for types with $c_i < K$, it is weakly dominant to bid the cap in a binding-cap environment. When the budget binds instead, bidders behave as in the no-cap case unless their no-cap bid would exceed K, in which case their bid is truncated at K. Types with $c_i \geq K$ are priced out and earn zero at any admissible bid.

For the case in which the cap is undisclosed, if we assume $K \sim H \equiv \mathrm{U}[0, \bar{K}]$ for some $\bar{K} > 0$, it follows that the optimal bidding strategy is given as follows.

[8] Formally, $\beta^{DNo}(c_i) = c_i + \frac{1}{1-F(c_i)} \int_{c_i}^{\bar{c}} [1 - F(x)]dx$, so bids exceed cost by a strictly positive mark-up that depends on the cost distribution.

Corollary 3. *Suppose a buyer's budget is non-binding within a discriminatory auction with an undisclosed bid cap such that $K \sim H \equiv U[0, \bar{K}]$ for some $\bar{K} > 0$. Then bidder i's equilibrium bid is given by,*

$$\beta^{DUnd}(c_i) = \frac{\bar{K} + c_i}{2}, \tag{6}$$

for $c_i < \bar{K}$.

In Corollary 3 the presence of a potentially binding bid cap (not budget) results in the equilibrium bidding strategy to bid the mid-point between the private cost and the maximum possible bid cap. Thus bidders will increase their bids above their costs but understand that the higher the bid the more likely their bid will be above the realized bid cap.

2.4.1 Comparison of bidders' mark-ups and design options

To focus on the effects of bid–cap design rather than the budget constraint itself, we consider the case of a non-binding budget. In the specific setting of $n = 2$ players with uniform costs, all three design options derived above yield optimal *Eff*, allowing a direct comparison of bidders' mark-ups and informational rents. To illustrate, consider those instances in which *Rent* is worst—where a sufficiently "low-cost" bidder i satisfies $c_i < K \leq \bar{K}$ (but $\bar{c} \geq \bar{m}$). Then we have the following comparison:

Table 1: Mark-up comparisons ($n = 2$)

Bid Cap Scenario	**Bidder i's Mark-up**
No Bid Cap	$\frac{\bar{c}-c_i}{2}$
Disclosed Bid Cap	$\frac{(2K-c_i)}{2}$
Undisclosed Bid Cap	$\frac{\bar{K}-c_i}{2}$

The *Rent* performance of each of the three auction mechanisms therefore comes down to the auctioneer's communication of the (known or unknown) bid cap distribution to the buyers. For example, suppose the auctioneer is contemplating imposing a disclosed cap K or informing the bidders there instead exists some undisclosed cap with upper bound $\bar{K}$. Then, with respect to this particular cost c_i bidder, Table 2 illustrates how the auctioneer will achieve the best *Rent* outcome.

Obviously, the auctioneer is unable to make an informed case-by-case decision in this sense since doing so relies on knowledge of bidder i's cost c_i. However, these cases illustrate the trade-off in *Rent* outcomes based on the choice of disclosed cap K or upper bound of the bid cap distribution, $\bar{K}$. However, while the above cases illustrate that *Rent* can be improved

Table 2: Auctioneer's design options

Auctioneer's best *Rent* design	**Condition**
No Bid Cap	if $\bar{c} < \min\{2K - c_i, \bar{K}\}$
Disclose a Bid Cap	if $K < \frac{1}{2}\min\{\bar{c} + c_i, \bar{K} + c_i\}$
Create an Undisclosed Bid Cap	otherwise

by either lowering K or $\bar{K}$, it is important that the auctioneer balance this with participation considerations: lowering either K or $\bar{K}$ too far excludes those with higher costs from entering.

2.5 Numerical Equilibrium Predictions for the Experiment Analysis

The complexity of our auction model for $n > 2$—particularly within the experimental setting—renders the traditional approach of solving the differential equation associated with profit maximization analytically intractable. The discrete nature of the allocation rules and the presence of various treatments (e.g., bid caps) introduce discontinuities and non-linearities that preclude closed-form solutions. Consequently, we employ computational methods to approximate the equilibrium bidding strategies over private independent costs drawn randomly from a uniform distribution $U[0, 10.00]$ with a disclosed bid cap of \$7.07 and an undisclosed bid cap random drawn from the same distribution. The simulations use $n = 4$ symmetric bidders, matching the experimental design with a total budget of $\bar{M} = 100$, implying an effective per-bidder budget of $\bar{m} = 25$ in the allocation rule.

We employ the Covariance Matrix Adaptation Evolution Strategy (CMA-ES) to numerically approximate the symmetric equilibrium bidding strategy in our auction model. CMA-ES is a stochastic, derivative-free optimization algorithm that iteratively updates a multivariate normal search distribution, adapting its covariance matrix to the topology of the objective function. This allows it to effectively learn and exploit the structure of complex, potentially non-smooth payoff functions without requiring gradient information (Hansen and Ostermeier, 2001)—an essential feature given the discontinuities introduced by budget and bid-cap constraints in our model. CMA-ES has been successfully applied in economics to solve dynamic stochastic general equilibrium models (Andreasen, 2010), fit behavioral game-theoretic models (Wright and Leyton-Brown, 2014), and address multi-agent resource allocation problems (Barbalios and Tzionas, 2014), underscoring its suitability for non-smooth objective functions like those arising in our auction context.

Our implementation of CMA-ES focuses on optimizing the parameters of the bidding strategy. We employ a general bidding function of the form $b(c) = c + ae^{-kc} - ae^{-10k}$, where a and

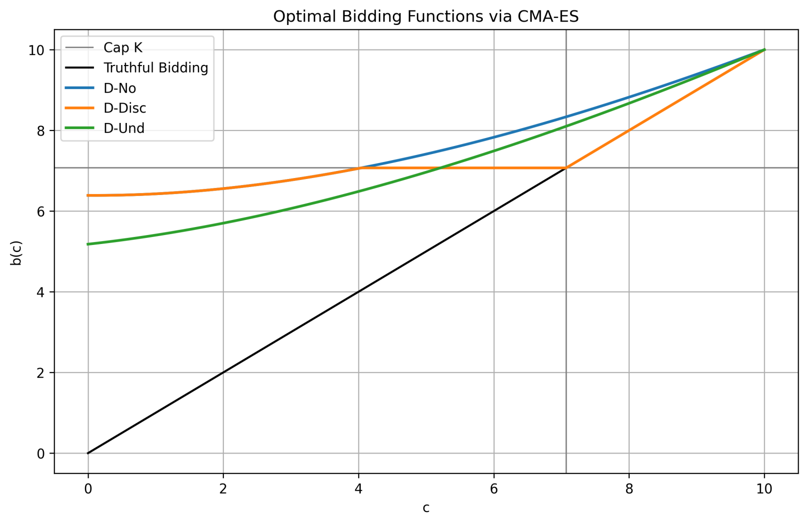

Figure 1: Estimated optimal bidding functions obtained via CMA–ES under three treatments: no bid cap (D–No), a disclosed bid cap (D–Disc), and an undisclosed bid cap (D–Und), with private costs drawn from $U[0, 10]$. Bids under an undisclosed cap lie strictly below those in the no-cap treatment for all cost types, reflecting additional precaution due to cap uncertainty. With a disclosed cap, bids coincide with no-cap bids for low costs but flatten at the cap as costs rise, leading to lower bids than in the no-cap case once the no-cap best response would exceed $K = 7.07$. At high cost levels, undisclosed-cap bids exceed disclosed-cap bids where the disclosed treatment imposes a hard ceiling.

k are parameters to be optimized. This functional form is selected for its desirable properties that align with theoretical expectations of bidding behavior in our auction context. Specifically, it ensures that $b(10) = 10$ (i.e., truthful bidding at the upper bound of the cost distribution) and that $b(c) \geq c$ for all c, which reflects it being a dominant strategy that bids should not be below costs. Moreover, $b(c) - c$ is decreasing in c, capturing the intuition that the markup over cost should diminish as costs increase. The estimated parameters associated with the optimal bidding strategy in the *DNo* and *DUnd* treatments are shown in Table 3.

Table 3: Estimated parameters for no-bid cap and undisclosed bid cap treatments

Treatment	a	k
DNo	10.11	0.10
DUnd	8.26	0.10

The optimal bidding strategy in the *DDisc* treatment is defined from Proposition 3. That is, the optimal bidding strategy follows $b(c) = c + ae^{-kc} - ae^{-10k}$ as defined in the *DNo* treatment for all $b(c) \leq K$, but then plateaus at $b(c) = K$, before following $b(c) = c$ for all $c \geq K$.

The estimated bidding functions are shown in Figure 1. The figure shows that optimal bids in the *DUnd* treatment are always strictly below those in *DNo* across the entire cost range except at $c = 10$, where they are equal. In contrast, optimal bids in *DDisc* equal those in *DNo* for low cost values ($c \leq 4.04$) but are lower in *DDisc* for medium and high cost values, so that, on average, bids are lower in *DDisc* compared to *DNo*. Thus, both bid cap treatments lower bids and, therefore, should lower rents compared to the *DNo* treatment. The comparison of optimal bids between the two bid cap treatments depends on the cost value: for low and medium costs ($c \leq 5.21$), the optimal bid in *DDisc* is higher, whereas for high cost values ($c > 5.21$) optimal bids in *DUnd* are larger.

In summary, our theoretical predictions can be grouped into the following bid- and market-level predictions:

Bid-level predictions.

- On average, both disclosed and non-disclosed caps reduce bids relative to a case with no bid cap.
- The relative ordering of bids under disclosed versus undisclosed caps depends on cost type: disclosed-cap bids are higher for low and medium costs, whereas undisclosed-cap bids are higher for high costs.
- Low-cost types follow the no-cap bidding strategy in *DDisc* (the cap does not bind), whereas medium- and high-cost types face truncation at K.
- In the undisclosed-cap treatment, precautionary shading leads to lower bids than in *DNo* even for low costs, since bidders internalise the risk of exceeding the unknown cap.

Market-level predictions.

- Informational rents are lower with either cap treatment compared to *DNo*.
- All treatments preserve (weak) allocative efficiency: bidding strategies remain monotone in costs in all three treatments, so the lowest-cost bidder always wins.

3 Experimental Design

3.1 Main Auction Task

The main task in the experiment is an auction task, which is designed to align closely with the theory developed in the previous section. Subjects take on the roles of bidders who are given a fixed capacity of five units of the good to sell, with each unit having the same cost of supply. Each subject is given a private independent cost for their goods drawn randomly from the uniform distribution $U[0, 10.00]$. Each subject only observes their own cost per unit at the start of each round, although the distribution from which the costs are drawn is common knowledge.

This environment allows us to measure bidding adjustments that cannot be identified in observational procurement data, such as focal-point bidding when the cap is disclosed or inadvertent overshooting when it is undisclosed. These behaviors arise from the distinct incentives created by each treatment: disclosure can shift bids upward for low- and medium-cost types by providing a salient reference point, whereas secrecy encourages precautionary shading but also exposes some bidders to the risk of exceeding the hidden cap. Because these forces operate in opposite directions and interact with the buyer's budget constraint, theory does not deliver a clear ranking of the two cap regimes. The experiment is therefore necessary to quantify the net effect of disclosure versus secrecy on rents and cost effectiveness.

Bidders earn a profit if they can sell a unit for a price greater than their cost per unit. Consistent with the theory, we restrict subjects to submitting a single price offer for all five of their units.[9] Bidders are placed into fixed groups of four. Each of these four bidders can offer to sell units to a single buyer, which in this experiment is the computer. The buyer is potentially constrained by their fixed budget of \$100, chosen by multiplying the total market capacity of 20 units by the mean of the private cost distribution \$5. To allow for learning, the auction task is repeated over 20 trading periods. These heterogeneous costs are redrawn in each of the 20 trading periods. While random—to enhance comparability across treatments—these costs were drawn prior to the experiment and the same sequences are used across all sessions and treatments.

3.2 Treatment Design

As shown in Table 4, we conducted three discriminatory auction treatments that vary the type of bid cap (either *None*, *Disclosed*, or *Undisclosed*). The first letter indicates the auction format (Discriminatory).

Table 4: Treatment Design and Number of Subjects

Treatment	**Subjects**
No Bid Cap (*DNo*)	*24*
Disclosed Bid Cap (*DDisc*)	*20*
Undisclosed Bid Cap (*DUnd*)	*20*

We set the bid cap equal to \$7.07, which we keep constant across rounds. As discussed earlier, choosing the bid cap level involves tradeoffs, yet the literature provides little guidance. Given our research question, we need a bid cap low enough to be binding sufficiently often but

[9] As mentioned above, it is a weakly (and, in some cases, strictly) dominant strategy—across all three treatments—to submit a price offer for all five units. In addition, having subjects only decide on their bid and not also the quantity to offer, reduces complexity and potential confusion in an already complicated auction design.

not so low that the implications are obvious. Balancing these considerations, we set the bid cap so that there is a 25% probability that all four bidders' private costs lie below it. In expectation, this implies the cap binds in 75% of rounds (i.e., at least one bidder's cost exceeds it). The resulting value of \$7.07 satisfies this calibration while avoiding a round-number focal point.

In the bid cap treatments, subjects are informed that any bids above the bid cap will be discarded (i.e., not considered) by the computer buyer. In the Disclosed treatment (*DDisc*), subjects are told the exact value of the bid cap, while in the Undisclosed treatment (*DUnd*) they are simply told that the bid cap will be selected from the same range of values as that of the private costs (i.e., on the interval $[0, 10]$), and that it will remain fixed for all trading rounds.[10] No mention of a potential cap is made in the No Bid Cap treatment (*DNo*).[11]

In all treatments, bidders are informed that, in each round, the computer buyer will purchase units in ascending price order until one of the following two conditions occurs: (1) the fixed budget of \$100 has been exhausted,[12] or (2) no more eligible bids remain.[13] This constitutes the competitive nature of the bidding process, in that bidders will not necessarily sell all five of their units in each auction round. Successful bidders receive a price equal to their offer value (bid) and earn a profit per unit equal to the difference between the price they received and their cost per unit, while unsuccessful bidders earn a profit of zero for the given auction round.

At the end of each auction round, each subject is informed of the number of units they sold (if any), the price at which the units were sold (if successful), and their total profit for the round. Subjects therefore learn only about their own result and not competitors. Feedback given is selected such that all subjects—irrespective of the treatment to which they are assigned—receive comparable information each round.[14]

[10] To prevent participants in *DUnd* learning the exact value of the bid cap from subjects in *DDisc*, we conducted all *DUnd* sessions first.

[11] Alternatively, to maintain parallelism we could have disclosed an arbitrarily high bid cap such as \$10 or even higher but were concerned that this might influence bids in an undesirable way.

[12] In the theoretical model, bidders either win or lose their entire unit. In the experimental implementation, the buyer may purchase a fractional remainder of the final bidder's units when the budget binds. This refinement ensures the budget is fully exhausted but does not materially alter strategic incentives: allowing infinitesimal divisibility simply smooths the indicator function in allocation decisions and yields the same equilibrium bidding function in the limit.

[13] Units are divisible to two decimal places. The tie breaking rule allocates an equal number of units to each bidder until the budget is exhausted.

[14] Our choice of minimal price feedback aligns with the practices implemented in the applied procurement auctions previously mentioned. For example, in the ERF, the Clean Energy Regulator only published the weighted average price paid across successful bids following each auction round (Commonwealth of Australia, 2014). This volume weighted price was heavily influenced by the size of projects contracted at each price point, and did not necessarily reflect the current price of the underlying good.

3.3 Implementation

In the experiment we use neutral language referring to each subject as a bidder of some unspecified underlying unit of a good. Full instructions for *DDisc* are provided in Appendix D.

The experiment proceeds as follows. The experimenter reads aloud the written instructions. Subjects then answer a series of 10-14 (depending on the treatment) quiz questions to check their understanding of the experiment and are paid A$0.50 for each correct response. The average mark recorded across these quizzes was 84%. After the 20 rounds of the main auction task are completed, subjects answer a short demographic questionnaire, followed by a scaled version of the Eckel and Grossman (2008) lottery choice task, to measure subject risk preferences. Each session lasted approximately 90 minutes, including the instruction time.

Subjects are paid their summed profits from three auction rounds randomly selected at the end of the experiment. We choose three rounds rather than one to lessen variation in payments and mitigate the earnings disadvantage faced by those subjects assigned a high cost in one of the selected rounds. Average subject earnings, including earnings from the quiz and risk tasks, were A$33.42 with an interquartile range of [$19.00, $43.50].[15]

Experimental subjects were recruited from the general pool of subjects at the University of Queensland using ORSEE (Greiner, 2015). All experiments were conducted using z-Tree (Fischbacher, 2007). Two-thirds of the subjects were undergraduate students, with half studying economics or business, and 30% male. Sample sizes are shown in Table 4, with any variation due to differences in no-show rates. Given that each auction comprises four bidders, we collected either five or six independent market observations per treatment, yielding a total of 16 independent observations at the aggregate market level.[16]

4 Discriminatory Auction Results

In this section, we evaluate the impact of bid caps in the discriminatory auction beginning at the aggregate market level, followed by individual bidding behavior.

4.1 Market Performance

In each market, we first examine two simple measures of the auction performance: the number of units purchased (*Units*), and the amount spent (*Spent*). However, because the number

[15] At the time of the experiment, this is equivalent to approximately US$23 with an IQR of US [$13, $29].

[16] Appendix Table C1 provides summary statistics of these demographic characteristics.

of units purchased and total spending may move together and do not fully capture allocative performance, we focus primarily on the three more comprehensive measures of auction performance defined in Definition 2.1.

As described above, we collected data from 16 unique market groups each participating in 20 rounds of auctions, thus yielding 320 market-level observations. From these we exclude 10 markets where the *Rent* ≤ 1 as these involve an average payment less than cost and result from bidder mistakes. We exclude two additional markets where the *Rent* > 10 as these extreme outliers could influence results. This leaves 308 market-level observations for our core market analysis.[17]

Table 5: Mean (S.D.) of Auction Performance Measures by Treatment

Treatment	n	Units	Spend ($)	BCE	Eff	Rent
DNo	119	14.9 (1.8)	100 (0)	6.8 (0.8)	4.4 (1.3)	1.7 (0.7)
DDisc	93	14.1 (3.2)	86 (18)	6.2 (0.5)	3.7 (1.2)	1.9 (0.8)
DUnd	96	13.0 (3.4)	80 (21)	6.2 (0.7)	3.6 (1.3)	2.0 (0.9)

Note: excludes markets where *Rent* ≤ 1 or *Rent* > 10.

Summary statistics are shown in Table 5 for our five measures of auction performance. We use Mann-Whitney tests of treatment differences to make some initial observations, with the unit of analysis being the group average, and reporting exact p-values from two-sided tests.[18] Consider first the simple measures of *Units* and *Spend*. Relative to no bid cap, an undisclosed bid cap significantly reduces both the number of units purchased ($p = 0.004$) and the amount spent ($p = 0.004$), while a disclosed bid cap has no impact on *Units* but does significantly reduce *Spend* ($p = 0.004$). Considering our main measure, *BCE*, both *DDisc* ($p = 0.004$) and *DUnd* (p=0.004) significantly improve (lower) *BCE* compared to *DNo*, but there is no difference between *DDisc* and *DUnd* ($p = 0.55$). Disaggregating the effect reveals that the improved *BCE* from a bid cap results from significant improvements in *Eff* in both *DDisc* ($p = 0.004$) and *DUnd* ($p = 0.004$) compared to *DNo*. On the other hand, *Rent* increases in *DDisc* and *DUnd* although the difference is only weakly significant in *DUnd* compared to *DNo* ($p = 0.08$).

Building on these initial observations, we next use random effects panel regressions to

[17]In total, we exclude only 3.8% of all discriminatory market-level observations. Appendix Table C2 reports exclusions by treatment and shows that markets with *Rent* ≤ 1 are more common in the bid cap treatments with three markets in *DUnd*, six in *DDisc*, and only one in *DNo*. These outcomes occur when one or more bidders mistakenly bid less than their cost. The other two excluded markets involve rents of 19.2 and 10.4, which the box-and-whisker plot in Appendix Figure C.1 shows to be outliers. Indeed, the next highest value of Rent is 5.9, while the mean is 1.9 and the median 1.5. These very high rents occurred when the only auction winners had very low costs (less than 0.7) but were paid prices ranging from 4-6.50 with the other costs (and bids) being sufficiently high.

[18]The top section of Appendix Table C3 reports p-values for all comparisons.

exploit our panel dataset, using the following specification:

$$Y_{jt} = \alpha + \beta_1 DDisc_j + \beta_2 DUnd_j + \gamma(1/Period)_t + u_j + e_{jt} \tag{7}$$

where j indexes the group (market) and t indexes the round. The dependent variable Y_{jt} measures either the *Units*, *BCE*, *Eff* or *Rent* of group j in round t, and $DDisc_j$ and $DUnd_j$ are time invariant treatment indicators (with DNo_j as the base treatment). We also include a trend variable $(1/Period)$ to control for learning, while u_j is the random effect and e_{jt} is the idiosyncratic error clustered at the matching group level. To analyze *Spend* we employ an equivalent probit regression where the dependent variable is an indicator of whether the full budget (of \$100) was spent or not (*Spend All*) to account for 63% of observations exhausting the budget. Coefficient estimates are reported in Table 6, where the bottom row of the table reports p-values from Wald tests comparing treatments where these are not obvious from the coefficients.[19]

From column (1) of Table 6 we see that significantly fewer units are purchased in *DUnd* compared to either *DDisc* ($p = 0.062$) or *DNo* ($p < 0.01$). However, column (2) shows that both types of bid cap make it significantly less likely that the full budget is spent ($p < 0.01$). From column (3), both bid-cap treatments yield significantly lower *BCE* values than the no cap baseline ($p < 0.01$), while the bottom panel shows no difference between the disclosed and undisclosed caps ($p = 0.875$). Column (4) reveals a similar pattern for *Eff*: both caps significantly improve efficiency relative to *DNo* ($p < 0.01$), with no distinction between cap types ($p = 0.518$). In contrast, column (5) shows that *Rent* increases significantly under the undisclosed cap ($p < 0.05$) but not under the disclosed cap, and the two cap treatments do not differ significantly from each other ($p = 0.448$). In Table 6 the trend term *(1/Period)* is significant for *BCE* and *Eff* but not for *Rent*, with the negative coefficients implying that both *BCE* and *Eff* increase (get worse) over time as subjects gain experience.[20]

These results imply that introducing a bid cap, whether secret or not, into a discriminatory auction improves *BCE* via an improvement in *Eff*—that is, lower cost units are purchased. However, there is a tradeoff, especially in the case of *DUnd*, as *Rent* significantly increases, though not by enough to fully offset the improvement in *Eff*. This contrasts with the *DDisc* case, where *Rent* increases but not significantly. Our results reveal that, contrary to expectations, bid

[19] In all cases our findings are consistent with those from the univariate tests reported above in the sense that all significant effects in the univariate tests are also found in the regressions, however, exploiting the panel data in the regressions and controlling for learning enables us to detect additional significant effects.

[20] These findings are robust to including data from all markets as reported in Appendix Table C4.

Table 6: Panel Regressions of Discriminatory Auction Performance

	(1) *Units*	(2) *Spend All*	(3) *BCE*	(4) *Eff*	(5) *Rent*
DDisc	-0.728	-6.161***	-0.655***	-0.685***	0.178
	(0.612)	(0.238)	(0.137)	(0.149)	(0.122)
DUnd	-1.944***	-6.450***	-0.626***	-0.782***	0.294**
	(0.368)	(0.226)	(0.171)	(0.128)	(0.136)
(1/Period)	2.085***	0.073	-0.974***	-0.469*	-0.165
	(0.501)	(0.455)	(0.162)	(0.256)	(0.125)
Constant	14.528***	6.030***	6.980***	4.484***	1.739***
	(0.227)	(0.160)	(0.096)	(0.097)	(0.075)
Observations	308	308	308	308	308
DDisc vs DUnd	0.062	0.220	0.875	0.518	0.448

Markets where *Rent* ≤ 1 or *Rent* > 10 are excluded from the analysis. The table reports estimated coefficients from linear regressions in columns (1) and (3)-(5) and from a probit regression in column (2).
Standard errors clustered on group (market) are in parentheses.
The omitted case is *DNo*. The bottom row reports p-values from treatment comparisons.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

caps do not in fact reduce rents. We summarize these findings in Result 1 below.

Result 1. *Implementing a bid cap in a discriminatory auction:*

a) Significantly improves budgetary cost-effectiveness regardless of whether the cap is disclosed or not. There is no difference in improvement if the cap is disclosed or not.

b) Significantly improves economic efficiency regardless of whether the cap is disclosed or not. There is no difference in improvement if the cap is disclosed or not.

c) Significantly increases information rents if the bid cap is undisclosed but has no impact if the bid cap is disclosed. There is no difference in improvement if the cap is disclosed or not.

4.2 Bidding Behavior

In this section, we examine individual bidding behavior, which together with the market institution contributes to the market performance results discussed in the previous section. Over the 20 rounds we have 1,280 individual bid observations. As subjects must enter a bid every round, our experiment was not designed to explicitly examine participation decisions. Nevertheless, subjects may enter a very high bid as a way of avoiding participation. To avoid such outlier values influencing our results, we exclude from our analysis the small number of bids greater than 15 (13 observations comprising just 1% of all bids), leaving 1,267 observations for our analysis.[21]

Figure 2 plots bids versus costs in each treatment with the black line showing the 45-degree line where *bid=cost*, and the gray horizontal and vertical lines at the bid cap level of 7.07. Consistent with the above, any bids over 15 are not shown. Recall the costs were randomly drawn from $U[0, 10.00]$ but kept the same across treatments to enhance comparability. Bidding below cost is a dominated strategy in all treatments regardless of cost. In *DNo* we observe very few bids below costs. Bidding below cost is also rare in *DDisc* and *DUnd* for costs below the bid cap value (around 1% of bids). However, for costs above the cap value we observe a large portion of bids below cost in the *DDisc* and *DUnd* treatments with more frequency in *DDisc*. While we explicitly alert subjects in the instructions that their profit may be negative if they submit an offer below their cost (even providing an example of this), we do not prevent them from doing so. We focus our analysis on rational bids only, i.e., bids that are greater than or equal to cost, excluding bids below cost from our analysis.[22] This leaves 1,186 bidding observations for our analysis in this section.

In Figure 2 the red lines plot the estimated optimal bidding functions (as derived in Section 2.5 and shown in Figure 1), while the green lines show the best linear fit of the data. In all treatments, bids increase with cost. From the theoretical predictions, in *DDisc* bids will be exactly at the cap if the budget is not binding; indeed, we observe a large concentration of bids just below the bid cap level. A similar pattern arises in *DUnd* although the data is noisier. Both *DDisc* and *DUnd* mechanically decrease the range of acceptable bids, however, as the figure shows, an undisclosed cap may result in some units with lower costs being excluded from participation. On the other hand, disclosing the cap seems to increase bids at the lower end of the cost range. These patterns illustrate the behavioral channel through which cap uncertainty can alter which bidders remain in the allocation set.

Table 7 reports summary statistics of bids for all cost values, as well as for three relevant ranges of costs, which we describe as 'low' (cost ≤ 4.04), 'medium' ($4.04 <$ cost ≤ 7.07) , and 'high' (cost > 7.07) costs, respectively.[23] Overall, bid caps tend to reduce both the average bid

[21] Appendix Table C5 reports exclusions by treatment, while the boxplot in Appendix Figure C.2 shows that bids above 15 are indeed extreme outlier values. (Note that for visibility, this figure excludes six bids over 85.) In *DNo* there are five bids above 15 with three made by one person. In *DDisc*, where bidding above the disclosed cap is clearly a way of choosing not to participate, there are eight bids > 31 made by only two people and no other bids exceeding 10. There are no such high bids in *DUnd*.

[22] Appendix Table C5 reports exclusions by treatment. Subjects may bid just to take part knowing the round is unlikely to count for earnings. In addition, these higher costs are typically inconsequential for auction outcomes. As seen in Table 5, average *Eff*, which reflects the average cost of purchased units, ranges from \$3.6 to \$4.4.

[23] While the intersection points in Figure 1 imply four potentially relevant ranges, in our analysis we limit this to only three relevant ranges due to the low number of observations with costs between 4.04 and 5.21, and for simplicity.

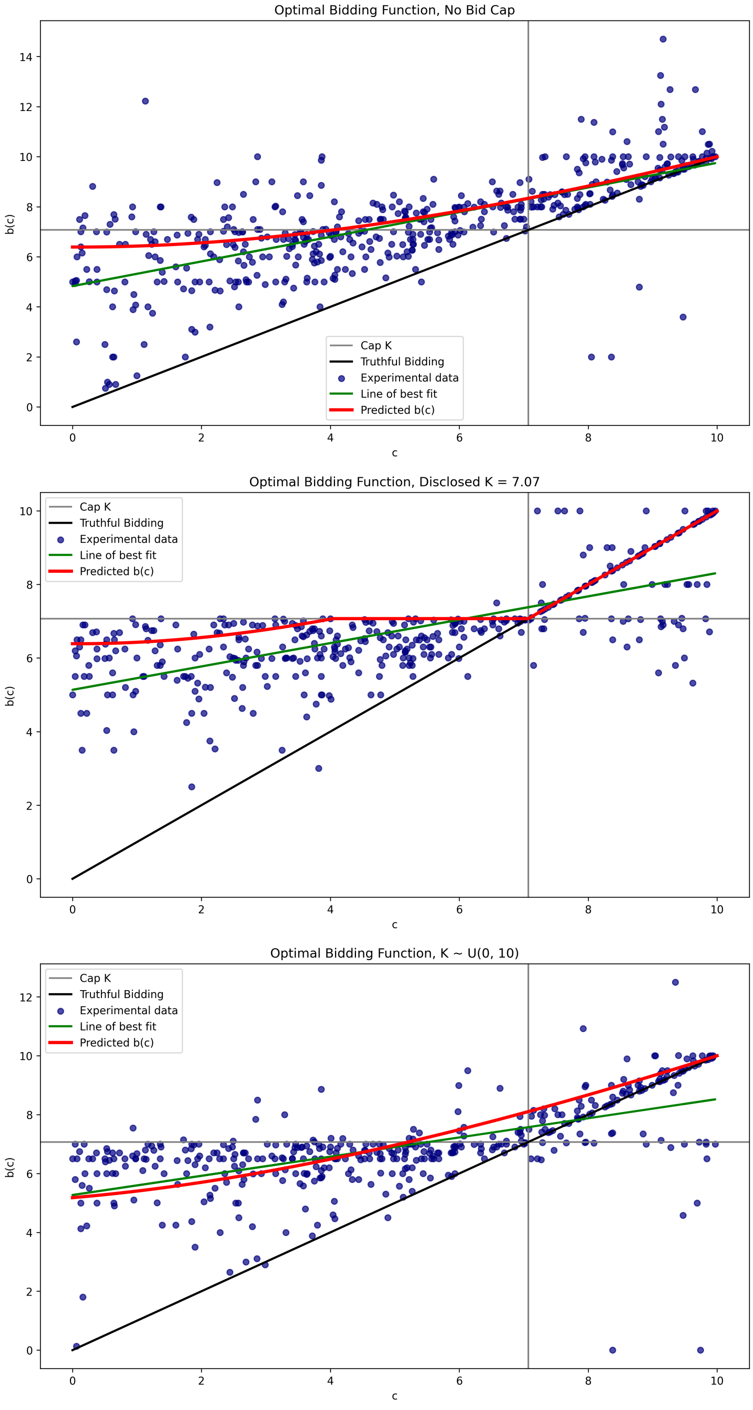

Figure 2: Optimal bidding strategy predictions and experimental data (bids over 15 are excluded).

and the variation in bids although the reduction in bids seems less apparent for low costs. Mann-Whitney tests of treatment comparisons demonstrate that bid caps of either type significantly reduce bids when costs are medium or high, and on average across all costs ($p < 0.02$ in all cases) but not when costs are low. When costs are medium, disclosing the bid cap decreases bids more than keeping it secret ($p = 0.02$), but no significant differences are detected across other cost ranges.[24]

As with the market level results, we estimate random effects panel regressions of bids using the following specification:

$$Bid_{ijt} = \delta + \mu_1 DDisc_j + \mu_2 DUnd_j + \theta Cost_{ijt} + \phi(1/Period)_t + u_i + e_{ijt} \tag{8}$$

where i indexes the individual, j the group, and t the round, $DDisc_j$ and $DUnd_j$ are time invariant treatment indicators (with DNo_j as the base treatment), $Cost_{ijt}$ is the cost, a Trend variable $(1/Period)$ to control for learning, u_i is the individual random effect and e_{ijt} is the idiosyncratic error clustered at the group level. We estimate the model including all cost values, and then separately for low, medium, and high costs (as defined above). The results are reported in Table 8.

Across all cost levels, imposing a bid cap significantly reduces bids whether the cap is disclosed ($p < 0.01$) or not ($p < 0.10$), with no difference between them. This difference between bids in *DNo* and *DUnd* is consistent with the optimal bidding functions plotted in Figure 1. Figure 1, however, also indicates that a proper comparison with bidding in *DDisc* requires segmenting the costs into the relevant ranges. When doing so, however, we observe no significant differences in bidding across any treatments when costs are low. For medium costs, we observe that bids with either type of cap are significantly lower than in *DNo* ($p < 0.05$), and that bidding in *DDisc* is significantly lower than *DUnd* ($p = 0.086$). This comparative bidding behavior approximates the optimal strategies that we observe in Figure 1. Finally, for high costs we observe that bids in both *DDisc* and *DUnd* remain significantly lower than in *DNo* ($p < 0.01$). Therefore, a bid cap significantly reduces bids except when costs are low, while disclosing the cap reduces bids for medium cost levels. These differences in bidding are consistent with improvement in BCE under bid caps reported in Result 1a and the difference in rents reported in Result 1c.[25] These differences are consistent with our findings in the previous

[24] Appendix Table C6 reports p-values for all comparisons.

[25] In Appendix Table C7 we show that these results are robust to controlling for bidders' individual demographic characteristics. We include controls for being male, locally born, age, studying business and/or economics, and being risk averse, with the latter defined as choosing one of the two most risk averse choices in the risk task. Of these demographics, only age and risk aversion significantly affect bids with both older subjects and risk averse

section . The results in Table 8 also confirm that bids are increasing in costs, and that significant learning occurs over the periods as bids tend to increase over time, except for when bidders have high costs. We summarize these results in Result 2 below.

Result 2. *Bidding behavior in discriminatory auctions*

a) On average across all costs, the presence of a bid cap whether disclosed or not significantly decreases bid levels.

b) Bid caps significantly reduce bids for medium and high cost levels, but not for low costs.

c) For medium costs, disclosed bid caps further dampen bids compared to undisclosed bid caps, while no differences in bids are detected for low and high cost values.

Table 7: Bidding Summary Statistics – Discriminatory Treatments

	All costs			Costs $\leq$ 4.04			4.04 $<$ Costs $\leq$ 7.07			Costs $>$ 7.07		
	Mean	SD	n	Mean	SD	n	Mean	SD	n	Mean	SD	n
DNo	7.4	1.9	468	6.1	1.7	190	7.2	0.8	143	9.4	1.2	135
DDisc	6.7	1.3	349	5.9	0.9	163	6.5	0.5	117	8.8	0.9	69
Dund	7.0	1.5	369	6.1	1.1	163	6.7	0.7	120	8.9	0.9	86

Note: Bids $>$ 15 and Bids $<$ Cost are excluded from the analysis.

Table 8: Panel Regressions of Bids in Discriminatory Auctions

	(1) All Costs	(2) Cost $\leq$ 4.04	(3) 4.04 $<$ Cost $\leq$ 7.07	(4) Cost $>$ 7.07
DDisc	-0.464***	-0.194	-0.664***	-0.540***
	(0.149)	(0.251)	(0.161)	(0.199)
DUnd	-0.269*	-0.050	-0.428**	-0.630***
	(0.153)	(0.309)	(0.198)	(0.191)
Cost	0.441***	0.201***	0.379***	0.816***
	(0.030)	(0.061)	(0.047)	(0.067)
1/Period	-0.570***	-1.167***	-0.238*	0.718**
	(0.169)	(0.217)	(0.142)	(0.310)
Constant	5.242***	5.837***	5.161***	2.367***
	(0.188)	(0.278)	(0.325)	(0.554)
Observations	1186	516	380	290
DDisc vs DUnd	0.315	0.646	0.086	0.419

Note: Bids $>$ 15 and Bids $<$ Cost are excluded from the analysis. Standard errors clustered on bidder are in parentheses. The omitted category is *DNo*. The bottom row reports p-values from treatment comparisons. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

subjects bidding lower amounts.

5 Extension: Bid Caps in a Uniform-Price Auction

In this section, we examine whether our experimental findings regarding bid caps are robust to a different auction format. A common alternative to a discriminatory auction in many applications is the uniform-price auction in which all successful bidders receive the same price (e.g., Latacz-Lohmann and Van der Hamsvoort, 1997). Considering alternative auction formats—such as the uniform price format—is important as it may result in distinct bidding behavior and market outcomes. Indeed, Balmford et al. (2023) compared the discriminatory versus uniform-price formats for Payments for Ecosystem Services (PES) and showed that mean bids reduced by 40% under the uniform-price rule compared to the discriminatory rule.

In a uniform-price auction all successful bidders receive the *same* 'market-clearing' price. The clearing price is equal to the minimum of the following: (i) the lowest bid offer at which the cost required to purchase the aggregate number of units *exceeds* the auctioneer's fixed budget, $\bar{M}$ (i.e., the lowest rejected bid); or (ii) the highest bid that is less than or equal to the cap. In the case of the uniform-price auction—where all successful bidders receive the same clearing price—we may define this $p_i^* = p^*$ more formally as:

$$p^*(\bar{M}, K) = \begin{cases} p & | \; S(p) \cdot p^* = \bar{M} \quad \text{if } S(p) \text{ flat} \\ \min p & | \; S(p) \cdot p^* \geq \bar{M} \quad \text{if } S(p) \text{ not flat} \\ \max p & | \; p^* \leq K \qquad\qquad \text{if } \bar{M} \text{ is non-binding} \end{cases} \tag{9}$$

We conducted additional experimental sessions using a uniform price auction format, with the experimental design identical to that used for the discriminatory auction sessions except for the clearing price determination, where now successful bidders are all paid the same price equal to the lowest rejected bid. Mirroring the discriminatory auction setup, we implemented three uniform-price auction treatments differing in the type of bid cap—*None*, *Disclosed*, or *Undisclosed*—as summarized in Table 9. The first letter indicates the auction format (Uniform Price).[26]

Table 9: Treatment Design and Number of Subjects

Treatment	**Subjects**
No Bid Cap (*UNo*)	*24*
Disclosed Bid Cap (*UDisc*)	*20*
Undisclosed Bid Cap (*UUnd*)	*24*

[26]As subjects were recruited from the same pool, their demographic characteristics and earnings were similar to those in the discriminatory treatments (refer to Appendix Table C1). While more male subjects participated in the uniform treatments, their lottery choices (a measure of risk preference) are almost identical.

5.1 Market Performance

Our analysis is conducted as in the previous section. We collected data from 17 unique market groups each participating in 20 rounds of auctions, thus yielding 340 market-level observations. From these we excluded 12 markets (3.5% of observations) where the *Rent* ≤ 1 and one market with *Rent* > 10 leaving 327 market-level observations for our core analysis.[27]

Summary statistics are shown in Table 10 for our five measures of auction performance. The number of *Units* purchased is significantly reduced in *UUnd* compared to both *UNo* ($p = 0.004$) and *UDisc* ($p = 0.03$). Correspondingly, *Spend* is also significantly lower in *UUnd* compared to both *UNo* ($p = 0.002$) and *UDisc* ($p = 0.009$). However, both *UDisc* ($p = 0.08$) and *UUnd* ($p = 0.04$) significantly improve *BCE* compared to *UNo*, with no difference between them. On the other hand, differences in *Eff* and *Rent* across treatments are small, and none are significant.[28]

Table 10: Mean (S.D.) of Auction Performance Measures by Treatment

Treatment	n	Units	Spend ($)	BCE	Eff	Rent
UNo	120	14.2 (2.2)	94 (10)	6.8 (1.5)	3.8 (1.4)	2.0 (0.8)
UDisc	93	14.2 (2.1)	92 (14)	6.5 (0.5)	3.9 (1.3)	2.0 (1.0)
UUnd	114	12.9 (3.0)	84 (21)	6.4 (0.7)	3.7 (1.4)	2.1 (1.2)

Note: excludes markets where *Rent* ≤ 1 or *Rent* > 10.

We next conduct random effects panel regressions using the specification in Equation 7, with results reported in Table 11. From columns (1) and (2) of Table 11 we see that significantly fewer units are purchased in *UUnd* compared to either *UDisc* ($p = 0.008$) or *UNo* ($p < 0.01$). However, both types of bid cap make it significantly less likely that the full budget is spent (*UDisc* vs. *UNo*, $p < 0.1$; *UUnd* vs. *UNo*, $p < 0.01$) with the impact significantly larger in *UUnd* compared to *UDisc* ($p = 0.05$). From column (3) we see that the coefficients for both *UDisc* and *UUnd* are significantly smaller than for *UNo* ($p < 0.05$) implying an improvement in *BCE*, while there there is no difference between *UDisc* and *UUnd* ($p = 0.891$). In columns (4) and (5) we observe no treatment differences for either *Eff* or *Rent*. Thus, similar to the discriminatory auction, we observe that both types of bid cap improve *BCE*, with no difference between them. In contrast, however, we find no impacts on the component parts of *Eff* and *Rent*. As in the discriminatory auction, we observe that both *BCE* and *Eff* deteriorate over time, as shown by the negative coefficients on the trend terms in columns (3) and (4) of Table 11. These findings

[27] As in the discriminatory auction, markets with *Rent* ≤ 1 are more frequent in the bid cap treatments with five markets in *UUnd*, seven in *UDisc*, and none in *UNo*. The box-and-whisker plot in Appendix Figure C.1 shows that the outcome with *Rent* > 10 (the actual value is 16.8) is indeed an outlier: the next highest value of *Rent* is 7, while the mean is 2.0, and the median 1.7. Appendix Table C2 summarizes these exclusions by treatment.

[28] As above, we use Mann-Whitney tests of treatment differences, with the unit of observation the group average, and reporting exact p-values from two-sided tests. See the middle section of Appendix Table C2 for p-values for all comparisons.

Table 11: Panel Regressions of Uniform-Price Auction Performance

	(1) Units	(2) Spend All	(3) BCE	(4) Eff	(5) Rent
UDisc	0.046	-0.311*	-0.365**	0.050	-0.051
	(0.492)	(0.169)	(0.158)	(0.183)	(0.120)
UUnd	-1.252***	-0.736***	-0.389**	-0.122	0.071
	(0.295)	(0.210)	(0.155)	(0.110)	(0.139)
(1/Period)	0.331	-0.569*	-0.723***	-0.826***	0.197
	(0.528)	(0.338)	(0.227)	(0.268)	(0.360)
Constant	14.138***	0.565***	6.965***	3.956***	1.998***
	(0.233)	(0.125)	(0.107)	(0.113)	(0.101)
Observations	327	327	327	327	327
UDisc vs UUnd	0.008	0.050	0.891	0.306	0.429

Markets where $Rent \leq 1$ or $Rent > 10$ are excluded from the analysis. The table reports estimated coefficients from linear regressions in columns (1) and (3)-(5) and from a probit regression in column (2). Standard errors clustered on group (auction) are in parentheses. The omitted case is UNo. The bottom row reports p-values from treatment comparisons.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

are summarized in Result 3 below.[29]

Result 3. *Implementing a bid cap in a uniform-price auction:*

a) Significantly improves budgetary cost-effectiveness regardless of whether the cap is disclosed or not. There is no difference in improvement if the cap is disclosed or not.

b) Has no significant impact on economic efficiency regardless of whether the cap is disclosed or not. There is no difference if the cap is disclosed or not.

c) Has no significant impact on information rents regardless of whether the cap is disclosed or not. There is no difference if the cap is disclosed or not.

5.2 Comparison of Auction Formats

In this subsection, we directly compare the performance of the two auction formats using a combined panel regression. The specification is as in equation 7 with added indicators for the three uniform treatments, and again with *DNo* as the base. As for our separate analysis of each format, we continue to exclude market observations with either $Rent \leq 1$ or $Rent > 10$. The results for our three main measures are reported in Table 12. From column (1) we observe no difference in *BCE* across the two formats in the absence of a bid cap (the coefficient on *UNo* is insignificant). However, from columns (2) and (3) we observe offsetting effects with

[29]Including data from all markets yields nearly identical results as reported in Appendix Table C8. The only minor change is that the difference in *Eff* between *UDisc* and *UUnd* becomes weakly significant ($p = 0.08$).

UNo significantly improving economic efficiency (decreasing the measure) but at the cost of increasing information rents. On the other hand, while *BCE* is significantly higher in *UDisc* compared to *DDisc* ($p = 0.050$) we observe no significant differences in the component parts. While *BCE* in *UUnd* is higher than *BCE* in *DUnd*, the difference is not significant ($p = 0.152$). Finally, the within-format comparisons are consistent with those reported earlier.[30]

Table 12: Panel Regressions of Comparative Auction Performance

	(1) BCE	(2) Eff	(3) Rent
DDisc	-0.654***	-0.683***	0.178
	(0.135)	(0.147)	(0.120)
DUnd	-0.629***	-0.781***	0.293**
	(0.168)	(0.127)	(0.134)
UNo	0.029	-0.591***	0.323***
	(0.127)	(0.128)	(0.099)
UDisc	-0.337**	-0.542***	0.273**
	(0.151)	(0.178)	(0.118)
UUnd	-0.360**	-0.713***	0.394***
	(0.148)	(0.104)	(0.136)
(1/Period)	-0.843***	-0.657***	0.025
	(0.141)	(0.186)	(0.198)
Constant	6.958***	4.517***	1.706***
	(0.092)	(0.092)	(0.078)
Observations	635	635	635
DDisc vs DUnd	0.884	0.510	0.447
UNo vs UDisc	0.019	0.784	0.666
UNo vs UUnd	0.011	0.259	0.604
UDisc vs UUnd	0.896	0.299	0.422
DDisc vs UDisc	0.050	0.466	0.488
DUnd vs UUnd	0.152	0.527	0.535

Markets where $Rent \leq 1$ or $Rent > 10$ are excluded from the analysis.
The table reports estimated coefficients from linear regressions.
Standard errors clustered on group are in parentheses.
The omitted case is *DNo.*
The bottom panel reports p-values from treatment comparisons.
$^{*}p < 0.10,^{**}p < 0.05,^{***}p < 0.01$

Result 4. *Comparison of Auction Formats*

a) In the absence of a bid cap, there is no significant difference in budgetary cost-effectiveness between the two auction formats. However, there is a trade-off as while information rents are significantly reduced in the discriminatory auction, economic efficiency is significantly worse.

b) With a disclosed bid cap, the discriminatory auction format significantly improves budgetary cost-

[30]Including data from all markets yields identical conclusions as reported in Appendix Table C9.

effectiveness compared to the uniform format.

c) *With an undisclosed bid cap, although the discriminatory auction improves budgetary cost-effectiveness relative to a uniform-price auction the difference is not significant.*

In contrast to Balmford et al. (2023)—who find a uniform-price format improves 'value-for-money' compared to a discriminatory auction—we find that there is no significant difference in BCE between formats in the absence of a bid cap. This is due to the information rents being significantly reduced under a discriminatory auction compared to the uniform-price auction. Although a uniform-price format significantly improves BCE compared to discriminatory format when the bid cap is disclosed the effect is not significant when the bid cap is undisclosed.

6 Discussion and Conclusions

This article investigates the (non)disclosure of bid caps in procurement auctions. We both theoretically and experimentally examine how the existence of a bid cap—and whether the cap is disclosed or undisclosed to participants—impacts bidding behavior and market outcomes. We first examine a discriminatory auction, where successful bidders obtain their bid, but later extend the model to incorporate a uniform-price auction in which all successful bidders obtain the same market clearing price. Our main focus is on budgetary cost effectiveness, which can be decomposed into two effects: an efficiency effect and information rent effect.

We find that adopting a cap—either disclosed or non-disclosed—significantly improves budget cost effectiveness. However, a non-disclosed cap significantly increases the information rents to participants. When we investigate the uniform-price auction, with either type of cap, we find budget cost effectiveness is improved, although there is no significant impact on either efficiency or information rents. Comparing a discriminatory versus uniform rule, we find that a discriminatory format significantly improves cost effectiveness when the cap is disclosed.

Bid caps are present in many contexts. Our analysis suggests that if policy makers want to implement a bid cap for cost effectiveness reasons, implementing either a disclosed or undisclosed cap will suffice. However, implementing an undisclosed cap appears to generate larger information rents. Overall, it is not a secret worth keeping.

References

Andreasen, M. M. (2010). 'How to maximize the likelihood function for a DSGE model', *Computational Economics*, **35**, pp. 127–154.

Arozamena, L. and Cantillon, E. (2004). 'Investment incentives in procurement auctions', *The Review of Economic Studies*, **71**, pp. 1–18.

Arozamena, L., Ganuza, J.-J. and Weinschelbaum, F. (2023). 'Renegotiation, discrimination and favoritism in symmetric procurement auctions', *European Economic Review*, **159**, p. 104566.

Balmford, B., Collins, J., Day, B., Lindsay, L. and Peacock, J. (2023). 'Pricing rules for PES auctions: Evidence from a natural experiment', *Journal of Environmental Economics and Management*, **122**, p. 102889.

Barbalios, N. and Tzionas, P. (2014). 'A robust approach for multi-agent natural resource allocation based on stochastic optimization algorithms', *Applied Soft Computing*, **18**, pp. 12–24.

Betz, R., Greiner, B., Schweitzer, S. and Seifert, S. (2017). 'Auction format and auction sequence in multi-item multi-unit auctions: An experimental study', *The Economic Journal*, **127**, pp. F351–F371.

Chassang, S., Kawai, K., Nakabayashi, J. and Ortner, J. (2022). 'Robust screens for noncompetitive bidding in procurement auctions', *Econometrica*, **90**, pp. 315–346.

Chassang, S. and Ortner, J. (2019). 'Collusion in auctions with constrained bids: Theory and evidence from public procurement', *Journal of Political Economy*, **127**, pp. 2269–2300.

Climate Change Authority (2017). 'Review of the Emissions Reduction Fund'.

Commonwealth of Australia (2014). 'Emissions Reduction Fund White Paper'.

Cramton, P., Hellerstein, D., Higgins, N., Iovanna, R., López-Vargas, K. and Wallander, S. (2021). 'Improving the cost-effectiveness of the Conservation Reserve Program: A laboratory study', *Journal of Environmental Economics and Management*, **108**, p. 102439.

De Silva, D. G., Kosmopoulou, G. and Lamarche, C. (2009). 'The effect of information on the bidding and survival of entrants in procurement auctions', *Journal of Public Economics*, **93**, pp. 56–72.

Dutra, J. and Menezes, F. (2005). 'Lessons from the electricity auctions in Brazil', *The Electricity Journal*, **18**, pp. 11–21.

Eckel, C. C. and Grossman, P. J. (2008). 'Forecasting risk attitudes: An experimental study using actual and forecast gamble choices', *Journal of Economic Behavior and Organization*, **68**, pp. 1–17.

Elyakime, B., Laffont, J. J., Loisel, P. and Vuong, Q. (1994). 'First-price sealed-bid auctions with secret reservation prices', *Annales d' Économie et de Statistique*, pp. 115–141.

Fischbacher, U. (2007). 'Z-tree: Zurich toolbox for ready-made economic experiments', *Experimental Economics*, **10**, pp. 171–178.

Friesen, L., Gangadharan, L., Khezr, P. and MacKenzie, I. A. (2022). 'Mind your Ps and Qs! Variable allowance supply in the US Regional Greenhouse Gas Initiative', *Journal of Environmental Economics and Management*, **112**, p. 102620.

Greiner, B. (2015). 'Subject pool recruitment procedures: organizing experiments with ORSEE', *Journal of the Economic Science Association*, **1**, pp. 114–125.

Hansen, N. and Ostermeier, A. (2001). 'Completely derandomized self-adaptation in evolution strategies', *Evolutionary computation*, **9**, pp. 159–195.

Hellerstein, D. (2017). 'The US conservation reserve program: The evolution of an enrollment mechanism', *Land Use Policy*, **63**, pp. 601–610.

Hellerstein, D. and Higgins, N. (2010). 'The effective use of limited information: Do bid maximums reduce procurement cost in asymmetric auctions?', *Agricultural and Resource Economics Review*, **39**, pp. 288 – 304.

Hyytinen, A., Lundberg, S. and Toivanen, O. (2018). 'Design of public procurement auctions: evidence from cleaning contracts', *The RAND Journal of Economics*, **49**, pp. 398–426.

Kagel, J. H. and Levin, D. (2002). 'Bidding in common-value auctions: A survey of experimental research', In *Common Value Auctions and the Winner's Curse*, Princeton: Princeton University Press, pp. 1–84.

Kagel, J. H. and Levin, D. (2016). 'Auctions: A survey of experimental research', In J. H. Kagel and A. E. Roth eds. *The Handbook of Experimental Economics*, **2**, Princeton: Princeton University Press, pp. 563–637.

Kawai, K. and Nakabayashi, J. (2022). 'Detecting large-scale collusion in procurement auctions', *Journal of Political Economy*, **130**, pp. 1364–1411.

Kirwan, B., Lubowski, R. N. and Roberts, M. J. (2005). 'How cost-effective are land retirement auctions? Estimating the difference between payments and willingness to accept in the Conservation Reserve Program', *American Journal of Agricultural Economics*, **87**, pp. 1239–1247.

Latacz-Lohmann, U. and Van der Hamsvoort, C. (1997). 'Auctioning conservation contracts: A theoretical analysis and an application', *American Journal of Agricultural Economics*, **79**, pp. 407–418.

Li, T. and Zheng, X. (2009). 'Entry and competition effects in first-price auctions: Theory and evidence from procurement auctions', *The Review of Economic Studies*, **76**, pp. 1397–1429.

Lopomo, G., Persico, N. and Villa, A. T. (2023). 'Optimal procurement with quality concerns', *American Economic Review*, **113**, p. 1505–29.

MacKenzie, I. A. (2022). 'The evolution of pollution auctions', *Review of Environmental Economics and Policy*, **16**, pp. 1–24.

Marion, J. (2007). 'Are bid preferences benign? the effect of small business subsidies in highway procurement auctions', *Journal of Public Economics*, **91**, pp. 1591–1624.

Nakabayashi, J. (2013). 'Small business set-asides in procurement auctions: An empirical analysis', *Journal of Public Economics*, **100**, pp. 28–44.

Porter, G. H., Conley, K. M., Gordon, C. V., Kneece Jr, R., Rieksts, B. Q., Shaw, A. H. and Tate, D. M. (2015). 'Affordability of defense acquisition programs'.

Pycia, M. and Woodward, K. (2025). 'A case for pay-as-bid auctions'. Available at SSRN: https://ssrn.com/abstract=5119595.

Riley, J. G. and Samuelson, W. F. (1981). 'Optimal auctions', *The American Economic Review*, **71**, pp. 381–392.

Rosenkranz, S. and Schmitz, P. W. (2007). 'Reserve prices in auctions as reference points', *Economic Journal*, **117**, pp. 637–653.

Rottner, C. (2025). 'Auctioning off budgets in procurement', *The RAND Journal of Economics, forthcoming*.

Schilizzi, S. G. (2017). 'An overview of laboratory research on conservation auctions', *Land Use Policy*, **63**, pp. 572–583.

Schilizzi, S. and Latacz-Lohmann, U. (2007). 'Assessing the performance of conservation auctions: an experimental study', *Land Economics*, **83**, pp. 497–515.

Tolmasquim, M. T., de Barros Correia, T., Addas Porto, N. and Kruger, W. (2021). 'Electricity market design and renewable energy auctions: The case of Brazil', *Energy Policy*, **158**, p. 112558.

Vincent, D. R. (1995). 'Bidding off the wall: Why reserve prices may be kept secret', *Journal of Economic Theory*, **65**, pp. 575–584.

Wright, J. R. and Leyton-Brown, K. (2014). 'Level-0 meta-models for predicting human behavior in games', In *Proceedings of the fifteenth ACM Conference on Economics and Computation*, pp. 857–874.

A Theoretical proofs for discriminatory auction format

Proof of Proposition 1. Let bidder i's cost per unit be c and suppose they submit a bid of $b(\tilde{c})$ per unit, where $\tilde{c} \in [\underline{c}, \bar{c}]$ is the "type" that bidder i mimics. Suppose the other $n-1$ bidders are following the equilibrium bidding strategy $\beta(c_j)$ for $j = 2, \ldots, n$, where $F(\cdot)$ and $f(\cdot)$ represent the common cumulative and probability density functions of these bidders' independent private costs, respectively. It is easily shown that it is a weakly dominant strategy that $b \geq c$.

If bidder i mimics type $\tilde{c}$, her expected payoff is,

$$\mathbb{E}[\pi_1] = (b(\tilde{c}) - c_i) \cdot \Pr(\text{bidder } i \text{ wins using } b(\tilde{c})). \tag{10}$$

The problem therefore reduces to computing bidder i's probability of winning a unit given an offer of $b(\tilde{c})$, the competing bids of the other $n-1$ bidders, and the auctioneer's fixed budget, $\bar{M}$. Given independence of the bidders' private costs, this probability is equal to,

$$\begin{aligned}
\Pr(\text{bidder } i \text{ wins using } b(\tilde{c})) &= \sum_{k=0}^{n-1} \left[\Pr\left(\left\{ b(\tilde{c}) + \sum_{j=1}^{k} b(c_{(j)}) \leq \bar{M} \right\} \cap \left\{ \tilde{c} = c_{(k+1)} \right\} \right) \right] \\
&= \sum_{k=0}^{n-1} \left[\Pr\left(\left\{ \sum_{j=1}^{k+1} b(c_{(j)}) \leq \bar{M} \right\} \right) F(\tilde{c})^k (1 - F(\tilde{c}))^{n-k-1} \right].
\end{aligned}$$

Substituting this probability term back into the payoff function (10) it follows that expected payoff is,

$$\mathbb{E}[\pi_i] = (b(\tilde{c}) - c_i) \cdot \sum_{k=0}^{n-1} G_k(\tilde{c}),$$

where,

$$G_k(\tilde{c}) := \Pr\left(\left\{ \sum_{j=1}^{k+1} b(c_{(j)}) \leq \bar{M} \right\} \right) F(\tilde{c})^k (1 - F(\tilde{c}))^{n-k-1}.$$

Differentiating this expression with respect to $\tilde{c}$ yields,

$$\frac{\partial \mathbb{E}[\pi_i]}{\partial \tilde{c}} = b'(\tilde{c}) \cdot \sum_{k=0}^{n-1} G_k(\tilde{c}) + (b(\tilde{c}) - c_i) \cdot \sum_{k=0}^{n-1} \left[f(\tilde{c}) G_k(\tilde{c}) \left\{ \frac{k - (n-1)F(\tilde{c})}{F(\tilde{c})(1 - F(\tilde{c}))} \right\} \right]. \tag{11}$$

The (symmetric) equilibrium can then be found by setting (11) equal to 0 and $\tilde{c} = c$. Rearranging, it follows that the equilibrium bidding strategy is implicitly defined by,

$$\beta'(c_i) = (\beta(c_i) - c_i) \frac{\sum_{k=1}^{n-1} f(c_i) G_k(c_i) \left\{ \frac{(n-1)F(c_i) - k}{F(c_i)(1 - F(c_i))} \right\}}{\sum_{k=0}^{n-1} G_k(c_i)}. \tag{12}$$

□

Proof of Proposition 2. We consider the case in which the buyer's budget is non-binding, so any admissible bid up to the cap would be accepted if $b \leq K$. Let c denote bidder i's per–unit cost and K the disclosed bid cap.

Case 1: $c \leq K$. Any admissible bid must satisfy $b \leq K$. Conditional on winning at bid $b \leq K$, the bidder's per–unit payoff is $b - c$, which is strictly increasing in b. Hence the payoff is maximized at the highest admissible bid, $b = K$. Any $b \neq K$ yields strictly lower expected payoff. Therefore $b^* = K$ is a *strictly dominant* strategy.

Case 2: $c > K$. Any admissible bid $b \leq K$ yields $b - c < 0$, while rejection yields 0. Thus any $b > K$ (guaranteeing rejection) yields a strictly higher payoff than any $b \leq K$. Therefore bidding $b > K$ is a *strictly dominant* strategy (equivalently, it is optimal not to participate). □

Proof of Proposition 3. We treat the two cost regions separately.

Case 1: $c_i \geq K$. Since bidding below cost is weakly dominated, bidder i must have $\beta(c_i) \geq c_i \geq K$. Any such bid is infeasible under the cap and is automatically rejected, yielding a payoff of zero. Because the bidder cannot earn a positive payoff at any admissible bid, any $\beta(c_i) \geq c_i$ is weakly optimal. In particular, bidding $\beta(c_i) = c_i$ is a weak best reply.

Case 2: $c_i < K$. We compare the disclosed-cap environment to the no-cap (D–No) environment.

(i) If $\beta^{DNo}(c_i) \leq K$: Then the cap never binds for type c_i. The strategic environment facing this type is identical to the D–No treatment (the cap is effectively irrelevant). Hence the unique symmetric equilibrium bid is $\beta^{DNo}(c_i)$; deviating would contradict optimality in the D–No game.

(ii) If $\beta^{DNo}(c_i) > K$: Since $c_i < K$, the bidder can earn strictly positive surplus by bidding at or below K. Any bid above K is strictly dominated because it is infeasible and yields zero payoff. Thus any best reply must satisfy $\beta(c_i) \leq K$.

In the D–No environment, the optimal bid would be $\beta^{DNo}(c_i) > K$. Given that the cap truncates the feasible set down to $[c_i, K]$, the bidder's optimal feasible action is the highest admissible bid. Therefore the best reply is $\beta(c_i) = K$.

Combining all cases establishes the stated bidding strategy. □

Proof of Proposition 4. Because the buyer's budget is non-binding and the auction is discriminatory, each bidder's payoff depends only on their own bid and cost. We can therefore analyze a representative bidder i in isolation.

Case 1: $c_i > \bar{K}$. If $\beta_i(c_i) \leq \bar{K}$, the bid would be accepted and yield payoff $\beta_i(c_i) - c_i \leq \bar{K} - c_i < 0$. If $\beta_i(c_i) > \bar{K}$, the bid is rejected and payoff is 0. Hence it is strictly dominant to bid at or above the cap, $\beta_i(c_i) \geq \bar{K}$.

Case 2: $c_i = \bar{K}$. Any accepted bid yields zero or negative profit, and rejection yields zero. Thus any $\beta_i(c_i) \geq \bar{K}$ is (weakly) optimal, and bids below $\bar{K}$ are strictly dominated.

Case 3: $c_i < \bar{K}$. Let $\beta_i(c_i) \equiv \beta_1$ denote bidder 1's bid. The bid cap K is a random variable distributed according to H. The bidder's payoff is

$$u_i = \begin{cases} 0, & \text{if } \beta_1 > K, \\ \beta_1 - c_i, & \text{if } \beta_1 \leq K. \end{cases}$$

Since the budget is non-binding, acceptance depends only on whether $\beta_1 \leq K$. The expected utility is therefore

$$\mathbb{E}[u_i(\beta_1)] = (1 - H(\beta_1))(\beta_1 - c_i),$$

which does not depend on other bidders' costs or strategies. The same reasoning thus extends immediately to $n \geq 2$ players.

Differentiating with respect to β_1 gives

$$\frac{d\mathbb{E}[u_i]}{d\beta_1} = -h(\beta_1)(\beta_1 - c_i) + 1 - H(\beta_1).$$

The value of β_1 that satisfies the first-order condition for an interior solution (denoted b) therefore satisfies,

$$\begin{aligned} 0 &= -h(b)(b - c_1) + 1 - H(b) \\ \Leftrightarrow \quad h(b)b &= h(b)c_1 + 1 - H(b) \\ \Leftrightarrow \quad b &= \frac{1 - H(b)}{h(b)} + c_1. \end{aligned} \tag{13}$$

The second derivative is

$$\frac{d^2\mathbb{E}[u_i]}{d\beta_1^2} = -h'(\beta_1)(\beta_1 - c_i) - 2h(\beta_1),$$

which is strictly negative if and only if,

$$-h'(\beta_1)(\beta_1 - c_1) - 2h(\beta_1) < 0 \quad \Leftrightarrow \quad h'(\beta_1)(\beta_1 - c_1) + 2h(\beta_1) > 0. \tag{14}$$

Hence, provided that the continuously differentiable function $h(\cdot)$ satisfies (14), the payoff–maximizing bid is

$$b \;=\; \min\left\{ \frac{1-H(b)}{h(b)} + c_1,\; \bar{K} \right\},$$

where the minimum enforces feasibility of the bid under the cap: if the unconstrained maximizer b° solving $b^\circ = \frac{1-H(b^\circ)}{h(b^\circ)} + c_1$ lies in $[0, \bar{K}]$ then $b = b^\circ$; otherwise, by concavity the objective is increasing on $[0, \bar{K}]$ up to b°, and the constrained maximizer is the corner $b = \bar{K}$. □

Proof of Corollary 1. Suppose there are two bidders. Let bidder 1's cost per unit be c and suppose they submit a bid of b per unit. Denote the other bidder's bid as $\beta(c_2)$. We assume $\beta(c_2)$ to be an increasing, and differentiable function, where $F(\cdot)$ and $f(\cdot)$ represent cumulative- and probability density functions of this bidder's private costs, respectively. It is easily shown that it is a weakly dominant strategy that $b \geq c$.

Denote $\bar{m} = \frac{\bar{M}}{2}$ as the amount of fixed budget per unit of the bidders' capacity. We argue that $b < \bar{m}$ is a strictly dominated strategy. Suppose not and both bidders submit offers strictly below $\bar{m}$. Then, $b + \beta(c_2) < 2\bar{m} < \bar{M}$ and both bidders are successful in their offers and sell their unit to the buyer. However, the buyer's budget is not exhausted. In the discriminatory price auction, $\bar{M} - b - \beta(c_2) = \varepsilon_D > 0$ remains. Hence, bidder 1 could increase their bid by a sufficiently small amount and receive a strictly higher payoff. Specifically, bidder 1 could submit an offer of $b + \varepsilon_D$ and receive $b + \varepsilon_D > b$. Thus, $b \geq \bar{m}$.

So far we have established that $b \geq \bar{m}$. However, to determine how much the optimal bidding strategy b is above $\bar{m}$, it is necessary to consider two cases. Firstly, suppose the budget $\bar{M}$ is sufficiently large such that $\bar{m} > \bar{c}$. That is, there is sufficient budget to ensure that even a bidder with the highest possible private cost of $\bar{c}$ can make a positive profit by submitting a competitive bid. In this case, we claim that the unique equilibrium bidding strategy is $b = \bar{m}$. Notice that the payoff under this strategy is $\bar{m} - c > 0$ for all $c \in [\underline{c}, \bar{c}]$. To show that bidder 1 has no incentive to deviate, it suffices to show that bidding $b > \bar{m}$ results in a stricter lower payoff (since we know $b \geq \bar{m}$). Indeed, suppose bidder 1 instead submits an offer of $b = \bar{m} + \varepsilon$ for some $\varepsilon > 0$. Then, since bidder 2 is following the strategy $\beta(c_2) = \bar{m}$, the total value of the submitted bids is $2\bar{m} + \varepsilon > \bar{M}$. But this means that the budget is insufficient to fund both offers. Hence, only bidder 2 is successful and bidder 1 receives a payoff of 0. A strictly worse outcome. Notice that if $c < \bar{m}$ for all c (which is guaranteed when $\bar{c} < \bar{m}$) then all bidders submit offers of $\bar{m}$.

Now suppose the buyer's budget is smaller such that $\bar{m} \in [\underline{c}, \bar{c}]$. If $c > \bar{M}$, then bidder 1 cannot profit by winning a unit and so should submit an uncompetitive offer (such as $b = c$). In the knife-edge case that $c = \bar{M}$, we assume that bidder 1 follows the weakly dominant strategy $b = c = \bar{M}$.

With these trivial cases established, let us now turn our attention to the more interesting case in which $c < \bar{M}$. That is, the case in which bidder 1 has the potential to earn strictly positive profits. We will show that the equilibrium bidding strategy is now strictly increasing in c and is equal to one's private cost plus some mark-up. The first thing to note is that if $c < \bar{M}$, then it is a weakly dominant strategy to bid $c < b \leq \bar{M}$ (since bidding above $\bar{M}$ effectively disqualifies bidder 1, yet there is a strictly positive probability of winning some $b - c > 0$ by bidding within the total budget $\bar{M}$).

Now, let us derive the expected payoff function. Since a successful bidder is always paid a price equal to their offer in a discriminatory price auction, it follows that bidder 1's expected payoff is,

$$E[\pi_1] = (b - c)\,\Pr(1 \text{ wins a unit}).$$

Let us focus on the value of $\Pr(1 \text{ wins a unit})$. There are two (mutually exclusive) cases in which 1 wins a unit:

1. Player 1 submits the strictly lower offer (i.e., $b < \beta(c_2)$).

2. Player 1 submits the strictly higher offer (i.e., $b > \beta(c_2)$), and the sum of both players' offer is less than or equal to the total budget $\bar{M}$.

Notice that in case 1 we also have the implicit assumption that bidder 1's (lower) offer is less than or equal to the total budget $\bar{M}$. However, this is already guaranteed by our above observation that $b \leq \bar{M}$ in the current context where $c < \bar{M}$. Notice that in the above delineation of cases, we also ignore the scenario in which both players submit identical offers as this would require that both players have identical private costs (since $b'(c) > 0$), which is a zero probability event.

But now, notice that we can also ignore case (2) above. Indeed, recall that all bidders follow $b \geq \bar{m}$. Hence, if $b > \beta(c_2) \geq \bar{m}$, then $b + \beta(c_2) > 2\bar{m} = \bar{M}$ and there is insufficient budget to sustain 1's offer. Hence, the probability that player 1 submits the strictly higher offer *and* the sum of both players' offer is less than or equal to the total budget $\bar{M}$ is zero.

Given these observations, it follows that bidder 1's expected payoff simplifies to,

$$E[\pi_1] = (b - c)\,\Pr(\{b < \beta(c_2)\}) = (b - c)(1 - F(\beta^{-1}(b)).$$

The first-order necessary condition therefore implies that the optimal bidding strategy must satisfy,

$$\begin{aligned}
\frac{\partial E[\pi_1]}{\partial b} &= (1 - F(\beta^{-1}(b)) - (b-c)f(\beta^{-1}(b))\frac{1}{\beta'(\beta^{-1}(b))} \\
0 &= (1 - F(c)) - (b-c)f(c)\frac{1}{b'(c)} \\
b'(c) &= (b-c)\frac{f(c)}{1-F(c)}, \qquad (15)
\end{aligned}$$

where in the second line we use the fact that $\beta^{-1}(b) = c$ at the symmetric equilibrium bidding strategy. To solve this differential equation, let us first re-write it as,

$$b'(c)(1 - F(c)) = (b-c)f(c). \qquad (16)$$

Next, observe by the product rule we have,

$$[b(c)(1-F(c))]' = b'(c)(1-F(c)) - b(c)f(c). \qquad (17)$$

But notice that the first term on the RHS of Equation (17) is equal to the LHS of Equation (16). Hence, substituting the expression in (16) into the RHS of (17) yields,

$$[b(c)(1-F(c))]' = (b-c)f(c) - bf(c) = -cf(c), \qquad (18)$$

which holds for all $c \in [\underline{c}, \bar{M}]$. Integrating both sides of (18) with respect to c it follows that,

$$\begin{aligned}
\int_c^{\bar{M}} [b(x)(1-F(x))]'dx &= \int_c^{\bar{M}} -xf(x)dx \\
b(\bar{M})(1-F(\bar{M})) - b(c)(1-F(c)) &= \int_c^{\bar{M}} -xf(x)dx \\
b(c)(1-F(c)) &= \bar{M}(1-F(\bar{M})) + \int_c^{\bar{M}} xf(x)dx \\
b(c) &= \frac{1}{1-F(c)}\left[\bar{M}(1-F(\bar{M})) + \int_c^{\bar{M}} xf(x)dx\right], \qquad (19)
\end{aligned}$$

where the third line follows our assumption that bidder 1 follows the weakly dominant strategy $b = c = \bar{M}$ in the knife-edge case that $c = \bar{M}$. Notice in the reasonable event that the *overall* budget is at least large enough to cover a single bidder at whatever their cost (i.e. $\bar{M} \geq \bar{c}$), we have $F(\bar{M}) = \Pr(c \leq \bar{M}) = 1)$, and so this optimal bidding reduces to,

$$b(c) \quad = \quad \frac{1}{1 - F(c)} \int_{c}^{\bar{c}} x f(x) dx. \tag{20}$$

We have now shown that the optimal bidding strategy is given by

$$\beta_i^{DNo}(c_i) = \begin{cases} \bar{m} & \bar{c} < \bar{m} \\ \frac{1}{1-F(c)} \int_{c}^{\bar{c}} x f(x) dx & \bar{m} \leq \bar{c} \leq \bar{M} \\ \frac{1}{1-F(c)} \left[\bar{M}(1 - F(\bar{M})) + \int_{c}^{\bar{M}} x f(x) dx \right] & \bar{M} < \bar{c} \\ c & \bar{M} \leq c. \end{cases}$$

□

Proof of Corollary 2. For the realistic scenario $\bar{M} \geq \bar{c} \geq \bar{m}$, using the expression derived in (20), it follows that,

$$\begin{aligned} b^{DNo}(c_i) &= \frac{1}{1 - \frac{c_i}{\bar{c}}} \left[\frac{1}{\bar{c}} \int_{c}^{\bar{c}} x dx \right] \\ &= \frac{1}{\bar{c} - c_i} \left[\frac{1}{2} x^2 \right]_{c_i}^{\bar{c}} \\ &= \frac{1}{2(\bar{c} - c_i)} \left[\bar{c_i}^2 - c_i^2 \right] \\ &= \frac{1}{2(\bar{c} - c_i)} (\bar{c} - c_i)(\bar{c} + c_i) \\ &= \frac{\bar{c} + c_i}{2} \end{aligned} \tag{21}$$

That is, the equilibrium bidding strategy is equal to cost plus some strictly positive mark-up. □

Proof of Corrollary 3. For the uniform case $H \sim \mathrm{U}[0, \bar{K}]$, we have $h(\beta_1) = 1/\bar{K}$ and $h'(\beta_1) = 0$, so

$$\frac{d^2 \mathbb{E}[u_i]}{d\beta_1^2} = -2h(\beta_1) < 0$$

for all $\beta_1 \in [0, \bar{K}]$. Hence the objective is strictly concave and the first-order condition yields the unique global maximum.

Solving (13) with the uniform distribution for H yields the solution,

$$\begin{aligned} b &= \frac{1 - b_i/\bar{K}}{1/\bar{K}} + c_i \\ \Leftrightarrow \quad b_i &= \bar{K} - b + c_i \\ \Leftrightarrow \quad 2b &= \bar{K} + c_i \\ \Leftrightarrow \quad b &= \frac{\bar{K} + c_i}{2}, \end{aligned} \tag{22}$$

where $b \in [0, \bar{K}]$ provided that $0 \leq c_i \leq \bar{K}$, which holds in the current case where $c_i < \bar{K}$. □

B Weakly dominant strategy to submit an offer to all units

Proposition 5. *It is a weakly dominant strategy to submit an offer for λ units.*

Proof.

No Bid Cap

Assume for simplicity that there are only 2 bidders.[31] Let $\beta(c_2)$ denote the offer submitted by the other bidder and let $\bar{m} = \bar{M}/2\lambda$ denote the amount of fixed budget per unit of the bidders' capacity. Now, suppose by way of contradiction that it is optimal to submit a bid, b, for $\xi < \lambda$ units. Then, bidder 1's expected payoff will be:

Table 13: Payoff Outcomes for Quantity Offer $\xi < \lambda$

Case	Condition	Payoff	Probability
1	$b < \beta(c_2)$	$\min\left\{\xi, \frac{\bar{M}}{b}\right\}[b - c]$	$1 - F(\beta^{-1}(b))$
2	$b > \beta(c_2), \beta(c_2) \leq \bar{m}$	$\min\left\{\xi, \frac{\bar{M} - \beta(c_2)\lambda}{b}\right\}[b - c]$	$[F(\beta^{-1}(b))][F(\beta^{-1}(\bar{m})]$
3	$b > \beta(c_2) > \bar{m}$	0	$[F(\beta^{-1}(b))][1 - F(\beta^{-1}(\bar{m})]$

where we assume that the other bidder is following the weakly dominant strategy by submitting a bid for λ units.

Now suppose that bidder 1 increases their bid by some arbitrary amount, $\epsilon > 0$, so that they now bid b for $\widetilde{\xi} \equiv \xi + \epsilon$ units. Then, bidding this quantity, their expected payoff would be:

[31] The result generalises to any choice of $n \geq 1$ bidders. Indeed, for $n > 2$ there are simply more cases to consider but the argument remains identical: by increasing one's offer from ξ to $\xi + \epsilon$, a bidder either strictly increases their expected payoff or leaves themselves no worse off.

Table 14: Payoff Outcomes for Quantity Offer $\widetilde{\xi} \equiv \xi + \epsilon$

Case	Condition	Payoff	Probability
1	$b < \beta(c_2)$	$\min\left\{\widetilde{\xi}, \frac{\bar{M}}{b}\right\}[b-c]$	$1 - F(\beta^{-1}(b))$
2	$b > \beta(c_2), \beta(c_2) \leq \bar{m}$	$\min\left\{\widetilde{\xi}, \frac{\bar{M}-\beta(c_2)\lambda}{b}\right\}[b-c]$	$[F(\beta^{-1}(b))][F(\beta^{-1}(\bar{m})]$
3	$b > \beta(c_2) > \bar{m}$	0	$[F(\beta^{-1}(b))][1 - F(\beta^{-1}(\bar{m})]$

Clearly then, the expected payoffs in rows 1 and 2 of Table 14 weakly dominant those in rows 1 and 2 of Table 13, since $\widetilde{\xi} = \xi + \epsilon > \xi \; \forall \, \epsilon > 0$. Since the expected payoffs in row 3 are equivalent in both tables, it holds that for all b, the expected payoff when submitting an offer for $\widetilde{\xi}$ units weakly dominates that when submitting an offer for ξ units. This contradicts $\xi < \lambda$ being the optimal quantity offer for bid, $b \geq c$. The same logic can be applied for all quantity offers, $\xi < \lambda$. Clearly, at a quantity offer of $\xi = \lambda$, this is the maximum expected payoff for a given $\beta(C_1)$, as bidder 1's capacity is capped at $\beta(C_1)$.

Note that in the event that the auctioneer's budget is sufficiently large such that it is non-binding (e.g., $\frac{\bar{M}-\beta(c_2)\lambda}{b} > \lambda$), it is a strictly dominant strategy to submit an offer for all λ units.

Disclosed or Undisclosed Bid Cap

Given a (disclosed or undisclosed) bid cap, K, there are two possible scenarios to consider:

1. $K > c$

2. $K \leq c$

Now, recall that it is a weakly dominant strategy to adhere to the rule $\beta(c_i) \geq c_i$ (regardless of whether the bid cap is disclosed or not to the bidders). Under the second scenario, the bidder therefore always earns a payoff of 0. This is true for any quantity offer. Hence, the result holds trivially.

Let us suppose then that the first scenario holds. Since the bidder submits an offer above (or equal to) c_i, they earn either (i) a strictly positive payoff on each unit bid for if the budget and bid cap are sufficiently large, or (ii) a payoff of zero otherwise. Hence, by the exact same argument as applied in the No Bid Cap treatment, it is weakly dominant to submit an offer for all λ units: bidding for any fewer units "leaves money on the table" if (i) holds, and provides no advantage if (ii) holds. $\square$

C Appendix: Additional Results

C.1 Figures

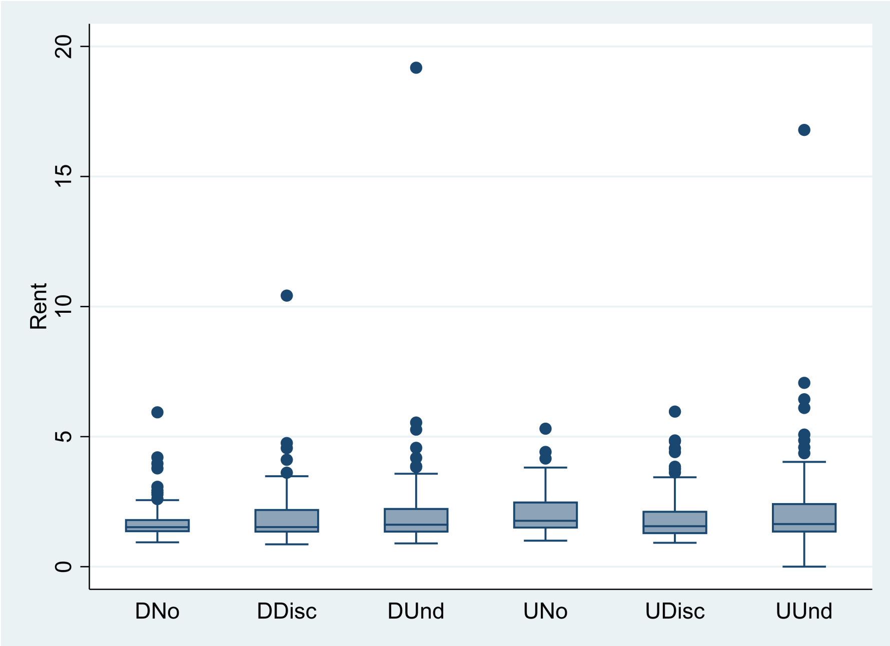

Figure C.1: Box-and-Whisker Plot of Rent Showing Outliers

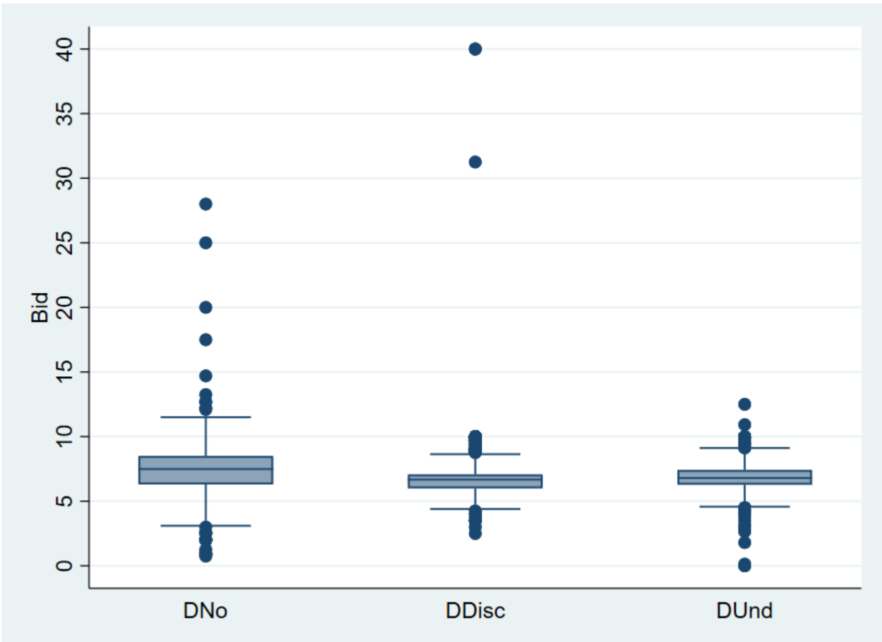

Figure C.2: Box-and-Whisker Plot of Bids Showing Outliers (excludes six observations with *bid*>85, one in *DNo* and five in *DDisc*).

C.2 Tables

Table C1: Mean (S.D.) of Demographic Measures

Label	Description	Discriminatory Treatments	Uniform Treatments
Male	Indicator of a male subject	0.30 (0.46)	0.49 (0.50)
Local Born	Indicator of a subject born in Australia or NZ	0.20 (0.41)	0.22 (0.42)
Age	Subject age in years	23.16 (7.17)	22.75 (4.44)
BusEcon	Indicator if subject is studying economics or business	0.50 (0.50)	0.62 (0.49)
Lottery	Lottery choice in risk task (between 1 and 5)	3.34 (1.46)	3.24 (1.48)
Risk Averse	Indicator for choosing one of the two most risk averse choices in the risk task	0.33 (0.47)	0.28 (0.45)
Quiz	% of quiz questions answered correctly	83.82 (15.62)	82.61 (12.50)
Earnings	Subject earnings (A$)	33.42 (18.98)	37.95 (21.12)

Table C2: Excluded Observations by Treatment

	Number of Markets		
	n	Rent$\leq$ 1	Rent$>$ 10
DNo	120	1	0
DDisc	100	6	1
DUnd	100	3	1
UNo	120	0	0
UDisc	100	7	0
UUnd	120	5	1
Total	660	22	3

Table C3: Exact p-values from two-sided Mann-Whitney tests of treatment differences, where the sign indicates if the first treatment is significantly larger (+) or smaller (-) than the second one.

Comparison	Units	Spend	BCE	Eff	Rent
DNo vs DDisc	0.43	+ 0.004	+ 0.004	+ 0.004	0.25
DNo vs DUnd	+ 0.004	+ 0.004	+ 0.004	+ 0.004	- 0.08
DDisc vs DUnd	0.22	0.15	0.55	0.69	0.42
UNo vs UDisc	0.79	0.43	+ 0.08	0.93	0.66
UNo vs UUnd	+ 0.004	+ 0.002	+ 0.04	0.48	0.82
UDisc vs UUnd	+ 0.03	+ 0.009	1.00	0.66	0.43
DNo vs UNo	+ 0.06	+ 0.002	1.00	+ 0.002	- 0.04
DDisc vs UDisc	1.00	0.22	0.22	0.69	0.55
DUnd vs UUnd	1.00	+ 0.05	0.33	0.43	0.66

The unit of observation is the average at the group/market level excluding markets with Rent$\leq$ 1 or Rent$>$ 10.

Table C5: Excluded Observations by Treatment

	Number of Bids			% of Bids<Cost (for bids$\leq$ 15)	
	n	> 85	$15 < bid < 85$	$Cost \leq 7.07$	$Cost > 7.07$
DNo	480	1	4	0.3	4
DDisc	400	5	3	1	36
DUnd	400	0	0	1	25
Total	1280	6	7	1	20

Note: Percentages of bids below cost are conditional on cost being below and exceeding the bid cap, respectively.

Table C4: Panel Regressions of Discriminatory Auction Performance – All Markets Included

	(1) Units	(2) Spend All	(3) BCE	(4) Eff	(5) Rent
DDisc	-0.827	-2.559***	-0.621***	-0.539**	0.216
	(0.616)	(0.358)	(0.130)	(0.220)	(0.180)
DUnd	-2.088***	-2.904***	-0.631***	-0.724***	0.438***
	(0.352)	(0.364)	(0.170)	(0.105)	(0.168)
1/Period	2.320***	-0.288	-1.056***	-0.424	-0.185
	(0.575)	(0.521)	(0.196)	(0.300)	(0.150)
Constant	14.515***	2.509***	6.984***	4.485***	1.737***
	(0.229)	(0.313)	(0.097)	(0.102)	(0.077)
Observations	320	320	320	320	320
DDisc vs DUnd	0.053	0.117	0.953	0.386	0.328

Notes: Columns (1) and (3)–(5) report linear regressions; column (2) reports a probit regression. Standard errors clustered by group (market) are in parentheses. The omitted baseline category is *DNo*. The bottom row reports p-values from treatment comparisons.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table C6: Exact p-values from two-sided Mann-Whitney tests of bidding in the discriminatory auction, where the sign indicates if the first treatment is significantly larger (+) or smaller (-) than the second one.

Comparison	**All Costs**	**Costs $\leq$ 4.04**	**4.04 < Costs $\leq$ 7.07**	**Costs > 7.07**
DNo vs DDisc	+0.001	0.41	+0.001	+0.02
DNo vs DUnd	+0.01	0.66	+0.02	+0.001
DDisc vs DUnd	-0.11	0.35	-0.02	0.49

The unit of observation is the average at the individual level
excluding any bid> 15 and any bid<cost.

Table C7: Panel Regressions of Bids in Discriminatory Auctions Controlling Demographics

	(1) All Costs	(2) Cost $\leq$ 4.04	(3) 4.04 $<$ Cost $\leq$ 7.07	(4) Cost $>$ 7.07
DDisc	-0.508***	-0.225	-0.745***	-0.555***
	(0.160)	(0.255)	(0.151)	(0.192)
DUnd	-0.286**	-0.049	-0.483***	-0.631***
	(0.123)	(0.270)	(0.158)	(0.170)
Cost	0.442***	0.201***	0.376***	0.815***
	(0.030)	(0.061)	(0.047)	(0.067)
1/Period	-0.571***	-1.171***	-0.222	0.720**
	(0.170)	(0.218)	(0.144)	(0.309)
Male	0.026	-0.121	0.156	0.168
	(0.136)	(0.267)	(0.115)	(0.138)
Born in Aus or NZ	0.037	0.133	0.004	0.156
	(0.191)	(0.316)	(0.114)	(0.247)
Age	-0.012***	-0.013*	-0.016***	-0.013*
	(0.004)	(0.007)	(0.005)	(0.008)
Bus./Econ. Student	0.121	0.107	0.150*	0.069
	(0.117)	(0.244)	(0.086)	(0.154)
Risk Averse	-0.177*	-0.167	-0.300***	-0.046
	(0.101)	(0.209)	(0.094)	(0.131)
Constant	5.526***	6.162***	5.562***	2.576***
	(0.226)	(0.315)	(0.321)	(0.505)
Observations	1186	516	380	290
DDisc vs DUnd	0.237	0.584	0.024	0.523

Note: Bids $>$ 15 and Bids $<$ Cost are excluded from the analysis. Standard errors clustered on bidder are in parentheses. The omitted category is *DNo*. The bottom row reports p-values from treatment comparisons. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table C8: Panel Regressions of Uniform-Price Auction Performance – All Markets Included

	(1) Units	(2) Spend All	(3) BCE	(4) Eff	(5) Rent
UDisc	0.048	-0.337*	-0.381**	0.228	-0.121
	(0.485)	(0.175)	(0.166)	(0.190)	(0.132)
UUnd	-1.446***	-0.749***	-0.543**	-0.098	0.129
	(0.241)	(0.224)	(0.222)	(0.129)	(0.227)
1/Period	0.283	-0.552	-0.759***	-0.917***	0.114
	(0.571)	(0.336)	(0.257)	(0.231)	(0.249)
Constant	14.146***	0.564***	6.972***	3.972***	2.013***
	(0.236)	(0.126)	(0.109)	(0.110)	(0.089)
Observations	340	340	340	340	340
UDisc vs UUnd	0.001	0.077	0.505	0.083	0.303

Notes: Columns (1) and (3)–(5) report linear regressions; column (2) reports a probit regression. Standard errors clustered by group (market) are in parentheses. The omitted baseline category is *UNo*. The bottom row reports p-values from treatment comparisons.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table C9: Panel Regressions of Comparative Auction Performance – All Markets Included

	(1) BCE	(2) Eff	(3) Rent
DDisc	-0.621***	-0.539**	0.216
	(0.128)	(0.217)	(0.177)
DUnd	-0.631***	-0.724***	0.438***
	(0.167)	(0.103)	(0.165)
UNo	0.041	-0.602***	0.330***
	(0.125)	(0.124)	(0.096)
UDisc	-0.340**	-0.374**	0.210
	(0.158)	(0.182)	(0.128)
UUnd	-0.502**	-0.700***	0.459**
	(0.214)	(0.121)	(0.223)
1/Period	-0.903***	-0.678***	-0.031
	(0.163)	(0.190)	(0.148)
Constant	6.956***	4.531***	1.709***
	(0.090)	(0.090)	(0.074)
Observations	660	660	660
DDisc vs DUnd	0.952	0.378	0.319
UNo vs UDisc	0.020	0.224	0.353
UNo vs UUnd	0.013	0.442	0.564
UDisc vs UUnd	0.499	0.078	0.296
DDisc vs UDisc	0.091	0.523	0.976
DUnd vs UUnd	0.597	0.824	0.934

Notes: All regressions are linear regressions. Standard errors clustered by group (market) are in parentheses. The omitted baseline category is *DNo*. The bottom rows report p-values from treatment comparisons.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

D Experimental Instructions

Below we provide the experimental instructions for the *DDisc* treatment. The instructions for the remaining five treatments follow an identical format and can be provided upon request.

Instructions

Introduction

Welcome to our experiment. Please read the following instructions carefully. The instructions are the same for everyone. If you follow these and make good decisions, you can earn money. To do so, you must remain quiet and not discuss any part of the experiment with other participants. Please turn off any electronic devices for the duration of the experiment. If you have any questions, raise your hand and we will come to you.

The experiment will consist of 20 rounds. An auction with four bidders (one of which is you) will take place in each round. All money values stated in the experiment are real money values. At the end of the experiment, all of the money you have earned will be paid to you in Australian dollars.

Auction Rules

You, and the three other participants, will bid in 20 rounds of an auction. Each round is a separate auction, and your bidding strategy in one round will not influence the outcome in subsequent rounds. This means that you can treat each round as a new auction.

Your Costs

In every round, each bidder will offer to <u>SELL</u> 5 units of a good to a single buyer, the computer. In each round, a bidder will have the same unit cost for all 5 units of the good. Think of this as being the dollar amount the bidder had to spend to produce *each* unit.

However, this unit cost may change from round to round. Also, the unit cost may be different for each of the four bidders.

Each bidder's unit cost will be shown to them at the start of each round, privately, on their computer screen. This means that your unit cost will not be shown to other bidders, and you will not see the unit costs of other bidders.

The figure below shows a screenshot of the cost screen privately revealed to each bidder at the start of each round.

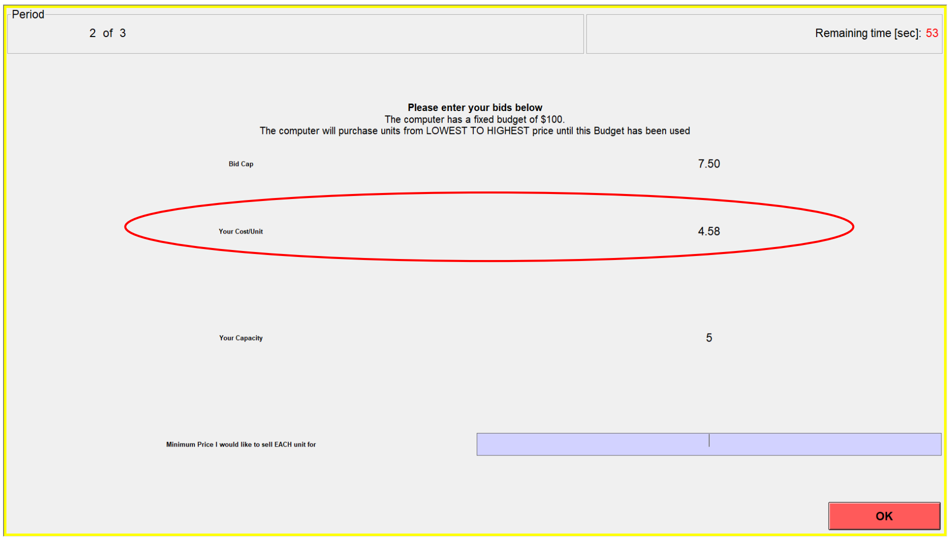

Each bidder's *unit* cost for the 5 units will be randomly drawn independently in each round from the interval:

$$0 \leq Unit\ Cost \leq 10$$

where the $Unit\ Cost$ is a real dollar amount to two decimal places. This means that every cent cost between $0, 0.01, 0.02, \dots, 9.98, 9.99, 10.00$ has an equal chance of being randomly assigned to a bidder in each round. It is important to remember that this is the cost per unit. For example, in the screenshot above, the bidder has a cost per unit of 4.58. This means that their *total cost* for *all 5 units* is $4.58 \times 5 = 22.90$.

Bidding Procedure

The design of the auction in each of the 20 rounds will remain as follows: each bidder will be shown their unit cost for each of their 5 units. Each bidder will then be asked to submit a price offer (per unit) on the computer. That is, each bidder must follow the following step:

1. Type the **minimum price** (to two decimal places) you are willing to sell each of your units for.

The figure below shows a screenshot of the entry each bidder must make at the start of each round.

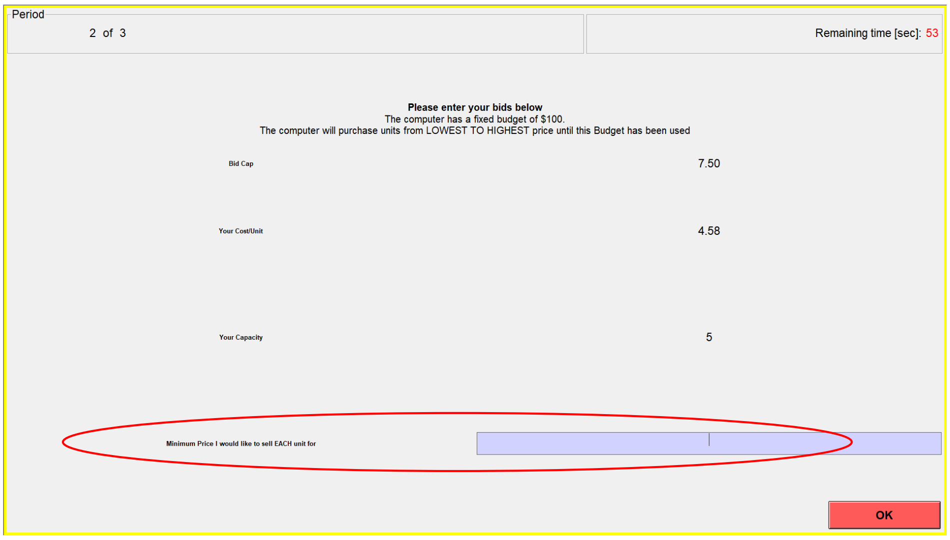

If the bidder is successful (this will be explained shortly), they will be paid a price per unit equal to the price indicated in Step 1 above (this will be explained shortly). The next round will operate identically, except that each bidder will be given a new unit cost, randomly drawn independently from the same interval, $0 \leq \text{Unit Cost} \leq 10$.

3

How Successful Bidders and Price are Determined

The bidding process is competitive. This means that each bidder may not actually sell all 5 units they submitted prices for in the bidding procedure. There are two factors that will affect who the successful bidders are, and how prices are determined:

1. The computer buyer has a fixed budget to spend in each round which cannot be exceeded.

1. There is a bid cap (per unit) on the maximum bid value accepted.
    - This means, that if you submit a price (per unit) <u>above</u> this bid cap, your offer will <u>not</u> be considered.
    - This bid cap will remain the <u>same</u> in all 20 rounds of the auction.

The bidders who successfully sell a unit(s) in each round will <u>*all*</u> be paid a price equal to the <u>bid price *they* submitted</u> (two or more bidders will receive the same price *only if* they submit bids with the same unit price). We explain how this price is calculated in the following example.

In the example below, we use a fixed budget of $80 and bid cap of $7.11 for illustrative purposes. These are **not** the same figures actually used in the auction you will participate in. The ***actual* figures that will be used in *your* experiment** will be given later.

EXAMPLE.

<u>Step 1: Order the bids from the four bidders from the lowest to the highest.</u>

Refer to the example in the table.

Bid	Number of Units Offered for Sale
5.80	5
6.70	5
7.10	5
10.90	5

<u>Step 2: Discard all bids above the bid cap.</u>

Suppose in this example, that a bid cap of $7.11 was selected by the computer before round 1. *The bid of $10.90 is discarded because it is above the bid cap.*

Bid	Number of Units Offered for Sale
5.80	5
6.70	5
7.10	5
~~10.90~~	~~5~~

Step 3: Calculate the total value of bids at each bid price.

This is computed as the 'Bid' × 'Number of Bids'. That is, Column 3 = Column 1 × Column 2.

Bid	Number of Bids	Value of Bids
5.80	5	5 x 5.80 = 29.00
6.70	5	5 x 6.70 = 33.50
7.10	5	5 x 7.10 = 35.50

Step 4: Starting from the first row, calculate the total cost (to the computer) of purchasing units as the bid price increases.

- ***Row one*** *is the* ***lowest bidder,*** *so the* ***total cost of bids*** *equals: $5.80 x 5 = $29.00;*
- *In* ***row two****, the total cost includes the 'Value of Bids' at $6.70* ***and*** *the 'Value of Bids' at $5.80, so the* ***total cost of bids*** *equals: $29.00 + $33.50 = $62.50.*
- ***Row three*** *is the* ***highest bidder****, so the total cost of bids is the* ***sum of all the 'Value of Bids' at each bid price****: $29.00 + $33.50 + 35.50 = $98.00.*

Bid	Number of Bids	Value of Bids	Total Cost to Computer
5.80	5	5 x 5.80 = 29.00	29.00
6.70	5	5 x 6.70 = 33.50	33.50 + 29.00 = 62.50
7.10	5	5 x 7.10 = 35.50	35.50 + 62.50 = 98.00

Step 5: Starting from the first row, determine whether the computer has enough budget to purchase the "Number of Bids" listed at each "Bid" price.

In this example, recall that the computer has a fixed budget of $80. We look at the 'Total Cost to Computer' column in each row and check whether the figure is less than or equal to $80. If it is, then the computer has enough budget, and we write "Yes"; otherwise, we write "No".

Bid	Number of Bids	Value of Bids	Total Cost to Computer	Enough Budget?
5.80	5	5 x 5.80 = 29.00	29.00	Yes
6.70	5	5 x 6.70 = 33.50	29.00 + 33.50 = 62.50	Yes
7.10	5	5 x 7.10 = 35.50	62.50 + 35.50 = 98.00	**No**

Step 6: Determine the highest accepted price.

The computer purchases units (up to two decimal places) in ascending price order until the fixed budget of $80 is exhausted. As can be seen in the table below, the computer has enough budget to purchase all of Bidder 1 and 2's offers, but ***not enough to purchase 5 units from Bidder 3 for 7.10.*** *That is, the computer's budget cannot afford the $98.00 recorded in the 'Total Cost to Computer'. The* ***highest successful bid is therefore $7.10*** *(though this bidder will only win a portion of their 5 units). All successful bidders are paid* ***different prices*** *(i.e., their own bid price).*

Bid	Number of Bids	Value of Bids	Total Cost to Computer	Enough Budget?
5.80	5	5 x 5.80 = 29.00	29.00	Yes
6.70	5	5 x 6.70 = 33.50	29.00 + 33.50 = 62.50	Yes
7.10	5	5 x 7.10 = 35.50	62.50 + 35.50 = 98.00	**No**

It is important to note that this price is calculated **after all bids above the bid cap have been discarded** (see Step 2). That is, this price will **always be less than or equal to the bid cap.**

Step 7: Determine the successful bidders.

All bidders who submitted a 'Bid' price ***below*** the highest accepted price (i.e., below \$7.10 in the example above) will be **paid their *own* bid price for all 5 of their units**. For example, in the table above, the bidder who submitted a 'Bid' for 5 units at a price of \$5.80 per unit would be paid $5 \times \$5.80 = \$29.00$, and the bidder who submitted a 'Bid' for 5 units at a price of \$6.70, would be paid $5 \times \$6.70 = \$33.50$. As indicated in the table above, this costs the computer \$62.50.

Since the computer has a total budget of \$80.00, there is \$17.50 remaining (i.e., \$80 – \$62.50). The computer then **purchases as many units as possible** (up to two decimal places) from the next highest bidder at the **highest accepted price** of \$7.10.

The number of units that can be purchased with remaining budget is given by the remaining budget ÷ final price, which in this example is:

$$\frac{\$17.50}{\$7.10} = 2.46$$

In this example, the bidder in row 3 of the above table would receive payment for these 2.46 units. Hence, he or she would be paid $2.46 \times \$7.10 = \$17.50$.

In the event of multiple bidders at the same price, the computer will purchase an equal number from each of the bidders who submitted bids at the highest successful bid price until the budget has been exhausted.

NOTE:

- In the above example, the computer's fixed budget of \$80 was **completely spent** in purchasing units from the bidders.
- However, the computer's fixed budget ***may* not always be exhausted.**
- This will happen if **enough bids have been discarded in Step 2** (i.e., there are enough bids above the bid cap)**,** such that the computer has **no more units available to buy** with the remaining budget.
- In the above example, this would have occurred **if the bid cap was \$6.00**, for example.
- In this case, **only the first bidder's offer at \$5.80 would have been accepted**, costing the computer only \$5.80 x 5 = \$29.00 (i.e., less than the budget of \$80).

How Your Profit is Determined

In your experiment today, the following **actual figures** are used:

- **Fixed Budget:** The computer has a fixed budget of \$100. This means that the computer can spend a **maximum of \$100 *each round.***
- **Bid Cap**: The bid cap is \$7.07 in every round. This means that if you submit a price (per unit) **above** \$7.07, your offer will **not** be considered.

It is important that you consider your price offer carefully. The price you submit will determine how much money you earn in that round. Suppose you are a successful bidder (as explained above). Then, you will make a profit on each unit that you sell for a price higher than your cost. Specifically,

$$\text{Profit per unit} = \text{Bid Price} - \text{Cost}$$

where the bid price is the amount you specify in Step 1 of the Bidding Procedure.
Hence, your total profit as a successful bidder in that round would be equal to,

$$\text{Total profit in round} = \text{profit per unit} \times \text{successful units won}$$

At the end of the 20 auctions rounds, you will be paid your total profits in THREE randomly selected rounds.

An example of profit calculation is as follows. Consider the **Example** illustrated in the previous section (see Step 1 in the How Successful Bidders and Price are Determined section):

Bid	Number of Bids
5.80	5
6.70	5
7.10	5
10.90	5

Recall that the highest successful bid price was \$7.10. Consider the bidder who submitted the bid offer of 5.80 in row 1. Since their bid price is below the highest successful bidder's, they would receive their bid price for all 5 units bid. That is, the bidder would receive $5 \times \$5.80 = \$29.00$. Now, suppose that bidder has a private cost of \$4.80 for each of their 5 units. Then, their profit per unit for this round would be:

$$\$5.80 - \$4.80 = \$1.00.$$

Since they are paid for all 5 units in their bid, their total profit for this round would be,

$$5 \times \$1.00 = \$5.00.$$

At the end of the 20 auctions, if this is one of the THREE rounds randomly selected for payment, the bidder will be paid \$5.00, plus their profit from the other randomly selected rounds.

It is worth pointing out two possible other outcomes in the above example.

1. The bidder in the final row, who submitted a bid price of \$10.90 for each of their 5 units.
    - Since this bidder submitted a price **above the bid cap** (\$7.11 in the example) in this round, their bid was **automatically discarded** from consideration.
    - Even if the computer's \$80 fixed budget had not been exhausted in fulfilling all lower bids, **this bidder would not be paid.**
    - Remember that this was the bid cap used in the example and **is not equal to the bid cap in your experiment.**

2. The bidder in row two, who submitted a bid price of \$6.70.
   - Suppose that this bidder has a **private cost of \$7.05** for each of their 5 units.
   - Since their bid price is below the highest successful bidder's, they would receive their bid price for all 5 units (i.e., receive $5 \times \$6.70 = \$33.50$).
   - However, since that bidder has a private cost greater than this final price they would make a loss. In particular, their payoff per unit would be,

     $$6.70 - 7.05 = -\$0.35.$$

     Since they are paid for all 5 units, their ***total loss*** for this round would be,

     $$5 \times \$0.35 = \$1.75.$$

   - Suppose that this round was randomly selected as one of the THREE rounds counting towards their final payment; this bidder would be paid their total earnings in the other randomly selected rounds, ***less* this amount**.

Auction Results

At the end of each auction round, you will be shown three numbers:

- The **number of units you sold** in the round;
- The **price you were paid** for each unit sold (if applicable);
- Your **total profit** for the round.

Only you will see these figures. This means that you will also not see the number of units and profits earned by other bidders.

A sample feedback screen is shown below.

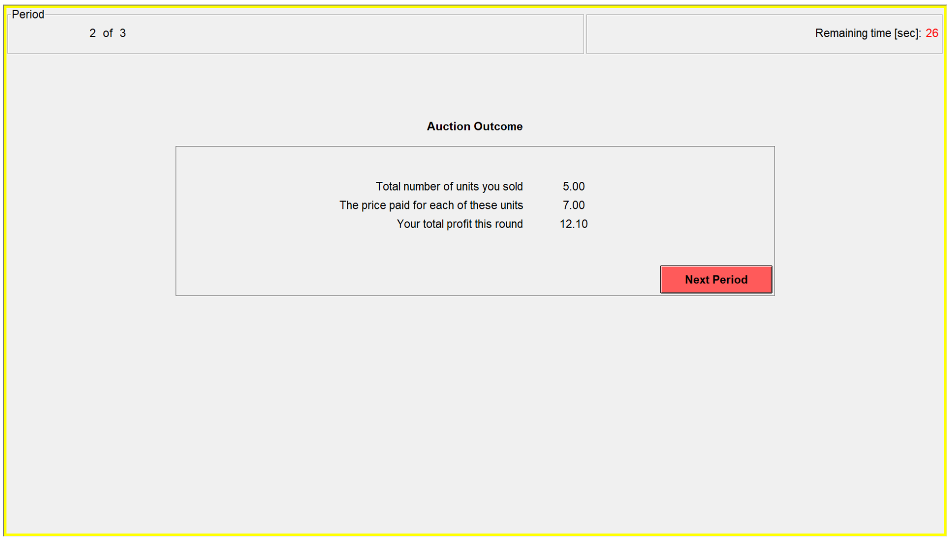

Auction Sets

After the auction results are displayed, the next round begins. There are **20 rounds in total**, and each follows an identical design. This means that the computer's fixed budget of $100 and the bid cap will be the same in all auctions. The other three bidders participating in each round will also remain the same. However, each round, you will be randomly given a new cost for each of your 5 units. This will be selected following the same process described earlier. The same will happen for the other three bidders.

When the 20th auction is complete, please wait quietly. We will then ask you to complete a short questionnaire. After you have completed the questionnaire, we will randomly select THREE of the 20 rounds and you will be paid the sum of your total profits in these three rounds.

Finally, note that each participant in your auction has been given the same rules. If you have any questions, please raise your hand and we will come to you. Please do not ask the other participants for help. Before starting the experiment, we will ask you 14 questions to check that you understand these instructions. You will earn $0.50 for each correctly answered question. Your total earnings from the experiment will be rounded up to the nearest 20 cents.

Summary

- You will be bidding in 20 auctions in a fixed group of four bidders.
- You will be bidding to sell your 5 units, where your cost for each of these units will be identical and randomly drawn from a uniform distribution between 0 and 10.
- This value will be randomly redrawn each round. The cost you receive is completely independent of other bidders' costs and your cost in any previous rounds.
- There is a bid cap of $7.07 per unit. All bids above this value will be discarded. The bid cap will remain the same for all 20 rounds.
- If you are a successful bidder, you will be paid the price you bid for each unit sold.
- The highest accepted bid price will be affected by other bidders' offers, the bid cap and the computer's fixed budget of $100 each round.
- If you submitted a bid price below the highest accepted bid price, you will always be paid the price you submitted for all 5 of your units.
- If you submitted a bid price equal to the highest accepted bid price, you will be paid the price you submitted for a proportion (and possibly all) of your 5 units.

- If you submitted a bid price above the highest accepted bid price, you will not be paid anything.

- For each unit you win, you will earn profit (or loss) = bid price – cost per unit.