Bidding in Multi-Unit Auctions under Limited Information

Bernhard Kasberger* and Kyle Woodward†

March 2025

Abstract

Multi-unit auctions frequently take place in environments with limited information, such as in new markets and under volatile macroeconomic conditions. We characterize optimal prior-free bids in such auctions; these bids minimize the maximal loss in expected utility resulting from uncertainty surrounding opponent behavior. We show that optimal bids are readily computable in this environment despite bidders having multi-dimensional private information. In the pay-as-bid auction the prior-free bid is unique; in the uniform-price auction the prior-free bid is unique if the bidder is allowed to determine the quantities for which they bid, as in many practical applications. We compare prior-free bids and auction outcomes across auction formats; while outcome comparisons are ambiguous, pay-as-bid auctions tend to generate greater revenue and welfare than uniform-price auctions when bidders' values are dispersed. We also compare outcomes in limited-information environments to outcomes in high-information environments, modeled as bidders playing Bayes-Nash equilibrium, and show that Bayes-Nash outcomes dominate prior-free outcomes when the auction is competitive.

JEL: D44, D81

Keywords: multi-unit auction, strategic uncertainty, robustness, regret minimization

*Department of Economics, University of Konstanz; bernhard.kasberger@uni-konstanz.de

†Assign Group; kyle@assign.group

We thank Justin Burkett, Piotr Dworczak, Marek Pycia, Orly Sade, Karl Schlag, Eran Shmaya, Alex Teytelboym, and Andy Zapechelnyuk, as well as audiences at INFORMS and the 2023 North American Summer Meeting of the Econometric Society for valuable feedback and comments.

1 Introduction

Pay-as-bid and uniform-price are multi-unit auction formats that play a critical role in the allocation of homogeneous goods. They are used to allocate generation capacity across power plants in electricity markets and to determine the interest rates at which governments can issue new debt.[1] In these auctions bidders submit demand curves to the auctioneer. The auctioneer uses the submitted demand curves to compute market-clearing prices and quantities, and each bidder is allocated their demand at the market-clearing price. In the pay-as-bid auction each bidder pays their bid for each unit received, while in the uniform-price auction they pay the constant market-clearing price for each unit received. The different pricing rules induce different strategic incentives and hence different outcomes. These differences are of interest to regulators: for example, in Summer 2022 prices for natural gas and electricity were high in Europe due to Russia's invasion of Ukraine. The high prices sparked a debate about whether electricity prices would be lower had the pay-as-bid auction been used instead of the uniform-price auction [Heller and Wieshammer, 2023].

Existing theoretical studies of these auctions typically analyze Bayes-Nash equilibria (BNE) and often restrict attention to relatively homogeneous bidders. A common justification for Nash equilibrium play is that players learn to mutually best respond over time. Indeed, Doraszelski et al. [2018] show that following a deregulation of the British electricity market, bidding behavior can be explained by Bayes-Nash equilibrium after three to four years. However, some multi-unit auctions happen rarely, close to a deregulation, or after substantial shocks that create structural and strategic uncertainty. For these auctions, Bayes-Nash equilibrium analysis is not applicable.[2] Moreover, when many goods are being auctioned, Bayesian models have previously been tractable only if the bidders were assumed to be homogeneous, while real-world bidders are often heterogeneous; little is known about equilibrium behavior in these auctions when bidders have general, multi-dimensional private values.[3]

[1] For government securities, see Brenner et al. [2009] and OECD [2021]. For electricity generation, see Maurer and Barroso [2011] and Del Río [2017].

[2] A short-run justification for equilibrium play is introspection: players reason to equilibrium strategies [Crawford, 2016]. In multi-unit auctions (as in any Bayesian game), this requires a common prior and commonly known equilibrium strategies [Aumann and Brandenburger, 1995]. However, even if bidders have a common prior, the computation of equilibrium strategies is typically intractable due to the multi-dimensionality of bidders' information [Swinkels, 2001; Hortaçsu and Kastl, 2012].

[3] Bayesian equilibrium constructions in these auctions do exist in parameterized contexts. For example, Engelbrecht-Wiggans and Kahn [2002] describe equilibrium when demand barely exceeds supply; Back and Zender [1993] and Wang and Zender [2002] when the good is divisible and bidders have common values; Ausubel et al. [2014] when bidders demand two units; Burkett and Woodward [2020a] when bidders' values are defined by order statistics; and Pycia and Woodward [2025] when bidders have common, decreasing marginal values.

In this paper we study initial play of arbitrarily heterogeneous bidders in multi-unit auctions.[4] In particular, we study the opposite extreme of Bayes-Nash equilibrium in terms of the bidders' information: While bidders in a BNE have common knowledge of the exact opponent bid distribution, our bidders know only that bids must be consistent with the rules of the auction; they bid under *maximal uncertainty*. We study bidders that deal with this uncertainty (ambiguity) by minimizing the maximal loss in expected utility due to not knowing the opponent bid distribution [Savage, 1951] and refer to their optimal bids as the *minimax-loss bids*.

A first takeaway is that our model is tractable for many questions that can be asked in multi-unit auction environments. For example, market outcomes can be computed where Bayes-Nash equilibrium methods are intractable (*cf.* footnote 2). In contrast with BNE, we characterize minimax-loss bids for any profile of weakly decreasing multi-dimensional bidder valuations.[5] As the optimal bid functions depend non-linearly on all marginal values, closed-form solutions are available only when the number of parameters is relatively low (such as in the case of two-unit demand or under flat marginal values); however, optimal bids can always be computed straightforwardly with numerical methods. Coming back to the motivating example of whether electricity prices would have been lower with an alternative pricing rule, one could specify the (distribution of) marginal costs of electricity providers and use our bid function characterizations to numerically investigate the impact of the market design on electricity prices. In terms of bidding language, our prior-free non-equilibrium approach is tractable with discrete and continuous bids and also in the empirically relevant setting where bidders are constrained to place fewer bids than the number of units they demand.[6]

A second finding is that minimax-loss bids under maximal uncertainty explain existing experimental bidding data qualitatively better than Bayes-Nash equilibrium. In experimental auctions with two-unit supply, two bidders, and flat marginal values, Engelmann and Grimm [2009] find that subjects do not play BNE strategies: They do not bid zero on the

[4] Level-k reasoning provides an alternative non-equilibrium approach that has been applied to initial play in multi-unit auctions [Hortaçsu et al., 2019]. However, Rasooly [2023] does not find support for the level-k model in an experiment designed to disentangle level-k from equilibrium behavior in single-unit auctions.

[5] In Appendix B we consider the possibility of complementarities in bidders' preferences, represented by increasing marginal values.

[6] In the constrained setting, the implied bid function is a step function, and the location and height of the steps are the bidders' choice variables. Although step functions are mathematically simple, they can be economically complex: when bids are constant over wide intervals, bidders are almost always rationed. When rationing occurs with positive probability, Bayesian equilibrium bids must take bidding incentives for non-local units into account, and the equilibrium first-order conditions imply a complicated non-local differential system [Kastl, 2012; Woodward, 2016]. Our prior-free approach is computationally more tractable, and we provide analytic solutions in the case of constant marginal values.

second unit in the uniform-price auction and they do not submit flat bids in the pay-as-bid auction. Such behavior can be qualitatively explained by minimax-loss bids under maximal uncertainty: the second minimax-loss bid is positive in the uniform-price auction and the second bid is lower than the first in the pay-as-bid auction.

Third, minimax-loss and Bayes-Nash equilibrium bids cannot be compared unambiguously. When there are two bidders and two units for sale in a uniform-price auction, the second bid is often zero in the Bayes-Nash equilibrium [Ausubel et al., 2014; Engelbrecht-Wiggans and Kahn, 1998; Noussair, 1995], but it is strictly positive under maximal uncertainty. In the pay-as-bid auction with single-dimensional private information and a uniform prior, the bids cannot be ranked unambiguously across informational regimes. However, when there are many goods and homogeneous bidders with sufficiently flat marginal values (as in Ausubel et al. [2014]), we show that, for relatively small quantities, Bayes-Nash equilibrium bids are higher than minimax-loss bids regardless of the auction format. The intuition is that common knowledge of homogeneity creates competitive pressure that leads to higher bids; under maximal uncertainty, bidders do not know that all are alike and that they will therefore bid similarly, which does not push bids up. On the other hand, for large quantities BNE bids in the pay-as-bid auction can be so high that they are above value [Pycia and Woodward, 2025]; that is, the equilibrium can be in dominated strategies. Minimax-loss bids are always undominated.

Regarding design implications, ex post payments are not generally comparable across auction formats.[7] For small quantities, the high bids of the uniform-price auction yield higher revenue than the low bids of the pay-as-bid auction, but for large quantities the low bids of the uniform-price auction yield lower revenue than the aggregate payment of both high and low bids in the pay-as-bid auction.[8] We also discuss how to set supply to maximize revenue (Propositions 5, 6, and 7).

Over time, auctioneers may be able to steer behavior toward the informational extremes of maximal uncertainty or Bayes-Nash equilibrium by revealing no or a lot of information about past bids (respectively). We find that the auctioneer does not unambiguously prefer one informational extreme over the other. When bidders are homogeneous, common knowledge of this homogeneity creates a highly competitive auction and revenue is higher in Bayes-Nash equilibrium than under maximal uncertainty, where bidders cannot rule out asymmetry. However, in settings where (tacitly) collusive, low-revenue Bayes-Nash equilibria exist—such as in a uniform-price auction with two bidders and two goods—revenue is higher under

[7]In an equilibrium framework with ambiguity and unit-demand bidders, Bougt et al. [2024] show that the pay-as-bid auction raises the highest revenue.

[8]Payment ambiguity has been observed in Bayes-Nash equilibrium [Ausubel et al., 2014] and empirically [Barbosa et al., 2022]

maximal uncertainty. In this case, bidders are unaware of the collusive scheme in which one bidder effectively offers a single good for free while pricing the second good prohibitively high. The extent to which this generalizes and whether intermediate informational policies can be optimal are beyond the scope of the paper.

Regarding empirical implications, our characterizations also lead to novel testable empirical predictions. For example, in the constrained setting and with constant marginal values, minimax-loss bid quantities are evenly spaced in the pay-as-bid auction, but concentrated on intermediate quantities in the uniform-price auction. In general, if one knew the bidders' values, then one could test whether they use the minimax-loss bids. Usually, however, one does not observe the bidders' values. In this case, the optimal bidding strategies can be used to estimate the bidders' values. Our uniqueness results and the characterizations of the optimal bids lead to point-identification of the values and a simple estimation procedure.

We introduce the model in the next section. Section 3 illustrates our approach and some findings with an analysis of the two-unit case. Section 4 contains some key theoretical results for the analysis of minimax loss in pay-as-bid and uniform-price auctions, which are applied in Sections 5 and 6 to analyze the unconstrained and bidpoint-constrained cases, respectively. Section 7 concludes. Omitted proofs are in Appendix A. Appendix B analyzes increasing marginal values. Appendices C and D contain a detailed analysis of the constrained case and the two-unit last accepted bid uniform-price auction, respectively.

2 Model

We consider an auction for quantity $Q > 0$ of a perfectly divisible, homogeneous good. There are $n \geq 2$ bidders participating in the auction. Buyer i, $i \in \{1, \ldots, n\}$, has marginal value function $v^i\colon [0, Q] \to \mathbb{R}_+$; that is, $v^i(q)$ is their marginal value for quantity q. We assume that marginal values are weakly decreasing, so that $v^i(q) \geq v^i(q')$ whenever $q \leq q'$, and assume as well that v^i is Lipschitz continuous.[9] For notational simplicity we assume that bidders have a strictly positive value for each unit, hence $v^i(Q) > 0$.[10] The marginal value functions $(v^i)_{i=1}^n$ may be distributed according to some joint distribution.

Bidder i submits a weakly decreasing and right-continuous bid function $b^i : [0, Q] \to \mathbb{R}_+$.

[9] We analyze increasing marginal values in Appendix B.

[10] Our results remain valid when bidders do not strictly demand all units, provided we replace aggregate supply Q with the supremum of all quantities for which marginal value is strictly positive, $\overline{Q}_i = \sup\{q\colon v^i(q) > 0\}$. Additionally, if $\overline{Q}_i < Q$, all results obtain in the limit with values $v^i(q) + \varepsilon$, letting $\varepsilon \searrow 0$.

After observing the bid profile $(b^j)_{j=1}^n$ the auctioneer computes a *market-clearing price* $p^\star$,

$$
\begin{aligned}
p^\star &\in \left\{p^{\text{LAB}}, p^{\text{FRB}}\right\};\\
p^{\text{LAB}} &= \sup\left\{p\colon \sum_{i=1}^n q_i \geq Q \text{ where } q_i = \sup\{q'\colon b^i(q') \geq p\}\right\},\\
p^{\text{FRB}} &= \inf\left\{p\colon \sum_{i=1}^n q_i \leq Q \text{ where } q_i = \sup\{q'\colon b^i(q') \geq p\}\right\}.
\end{aligned}
$$

There is strict excess supply for prices above p^{LAB} and strict excess demand for prices below p^{FRB}. Hence, the prices p^{LAB} and p^{FRB} are related to the last bid accepted and the first bid rejected, respectively.[11] All bids strictly above the market-clearing price $p^\star$ are awarded, and all bids strictly below the market-clearing price are rejected. When there are multiple bids placed at the market-clearing price, ties are broken randomly.[12]

Bidders are risk neutral. If a bidder with value v^i receives quantity q_i and makes transfer t_i, their utility is

$$\hat{u}\left(q_i, t_i; v^i\right) = \int_0^{q_i} v^i(x)\,dx - t_i.$$

We consider two common auction formats. In a *pay-as-bid* (or *discriminatory*) auction (PAB), transfers are equal to the sum of bids for received units, $t_i^{\text{PAB}} = \int_0^{q_i} b^i(x)dx$. In a *uniform-price* auction (UPA), transfers are equal to the market-clearing price times the number of units received, $t_i^{\text{UPA}} = p^\star q_i$. We analyze uniform-price auctions with p^{FRB} or p^{LAB} as the market-clearing price. The exact market-clearing price matters only when selling discrete units (as in Section 3). If opponent bids b^{-i} are distributed according to the integrable distribution B^{-i}, we write bidder i's interim utility as $u(b^i, B^{-i}; v^i) = \mathbb{E}_{B^{-i}}[\hat{u}(q^i(b), t^i(b); v^i)]$, where q^i and t^i are functions that map, according to the auction rules, the bidders' bid functions $b = (b^i, b^{-i})$ to bidder i's quantity q_i and transfer t_i, respectively.[13]

[11] See Burkett and Woodward [2020a]. Treasury auctions frequently apply last-accepted-bid pricing (e.g., the United States and Switzerland) while theoretical analyses frequently study first-rejected-bid pricing [Ausubel et al., 2014].

[12] As long as all bids strictly above the market-clearing price are awarded, the precise tiebreaking rule does not affect our results.

[13] Integrability of B^{-i} is not a constraint on our results, since in all auction formats $\hat{u}$ is bounded below by 0 and above by $Qv^i(0)$.

2.1 Loss and regret

Given a distribution of opponent bids B^{-i}, bidder i's *loss* from bidding b^i instead of the interim-optimal bid is

$$L\left(b^i; B^{-i}, v^i\right) = \sup_{\tilde{b}} \mathbb{E}_{B^{-i}}\left[\hat{u}\left(q^i(\tilde{b}, b^{-i}), t^i\left(\tilde{b}, b^{-i}\right); v^i\right) - \hat{u}\left(q^i\left(b^i, b^{-i}\right), t^i\left(b^i, b^{-i}\right); v^i\right)\right].$$

Loss measures the difference between expected utility given bid b^i and the utility obtainable by optimizing the submitted bid with respect to distribution B^{-i}. For example, when bid b^i is a best response to distribution B^{-i}, loss is zero. Loss is evaluated from an interim perspective; the equivalent ex post concept is *regret*,

$$R\left(b^i; b^{-i}, v^i\right) = \sup_{\tilde{b}} \hat{u}\left(q^i\left(\tilde{b}, b^{-i}\right), t^i\left(\tilde{b}, b^{-i}\right); v^i\right) - \hat{u}\left(q^i\left(b^i, b^{-i}\right), t^i\left(b^i, b^{-i}\right); v^i\right).$$

Regret measures how much additional utility the bidder could receive if they had known the bids their opponents submitted prior to choosing their own bid.[14] A utility-maximizing bidder with perfect foreknowledge of their opponents' bids will have zero regret.

If bidder i knew the true distribution of opponent bids B^{-i}, she would evaluate potential bids by standard expected utility. However, in our model bidders face ambiguity regarding the true distribution B^{-i} and know only that $B^{-i} \in \mathcal{B}$, where $\mathcal{B}$ is a set of feasible distributions over opponent bids. In the presence of this ambiguity, bidder i evaluates potential bids according to the maximum loss generated by any feasible distribution of opponent bids; the optimal bid $b^\star$ minimizes this loss:

$$b^\star \in \arg\min_{b^i} \sup_{B^{-i} \in \mathcal{B}} L\left(b^i; B^{-i}, v^i\right).$$ [15]

Hence, our bidders adopt the same interim perspective as Bayesian players that best respond to (their belief of) the opponent bid distribution.[16]

We refer to $b^\star$ as bidder i's *minimax-loss* or *optimal* bid. We focus on the case of *maximal uncertainty*, in which $\mathcal{B}$ contains all joint distributions on feasible bid functions; i.e., all distributions over $n-1$ weakly-decreasing functions mapping $[0, Q]$ to $\mathbb{R}_+$. Note that $\mathcal{B}$ is rich enough to include uncertainty about the number of bidders and supply.[17]

[14] We follow Schlag and Zapechelnyuk [2021] and Kasberger and Schlag [2024] and refer to the interim concept as loss and to the ex post equivalent as regret.

[15] We show that a bid that minimizes worst-case loss always exists.

[16] The ex post perspective is frequently applied to the analysis of environments with ambiguity [Stoye, 2011, Bergemann and Schlag, 2011]. Our interim perspective is consistent with standard Bayesian analysis.

[17] There are bid distributions in $\mathcal{B}$ that put all the mass on bidder j bidding zero, i.e., $b^j(q) = 0$ for all q. This effectively reduces the number of bidders so that n is merely an upper bound on the number of bidders.

Savage [1951] introduced the minimax loss (regret) decision criterion for statistical decision problems. Since then it has been applied in econometrics [Manski, 2021], mechanism design [Bergemann and Schlag, 2008, 2011; Guo and Shmaya, 2023, 2025], operations research [Perakis and Roels, 2008; Besbes and Zeevi, 2011], and more generally in strategic settings. Our paper belongs to the latter category. A first paper on analyzing games with minimax regret as the players' decision criterion was Linhart and Radner [1989] who study the minimization of worst-case regret in bargaining. Parakhonyak and Sobolev [2015] consider Bayesian firms best responding to consumers whose search rules for the lowest price are derived from worst-case regret minimization. Renou and Schlag [2010], Halpern and Pass [2012], Kasberger [2022], and Schlag and Zapechelnyuk [2023] propose solution concepts for loss (regret) minimizing players.

We offer a descriptive and a prescriptive interpretation of minimax-loss bids. From a prescriptive perspective, a practical advantage of our non-Bayesian approach is that the bids are completely prior-free, i.e., they do not depend on the other bidders' value distributions and strategies. All a bidder needs to know is their own willingness-to-pay. The bids are robust because the bidder need not worry about misspecified beliefs. Indeed, if any bid distribution is deemed possible, then in particular the actual distribution is possible. Kasberger and Schlag [2024] illustrate empirically that loss-minimizing bids perform well in first-price auctions, despite bidders having very coarse beliefs about competitors' behavior. Even the bid under maximal uncertainty performs better than observed bids in experimental and field data. On the other hand, group decision-making provides a descriptive motivation for minimax loss. Suppose a corporation tasks a team with finding the right bid. Based on information learned after the auction, the executive board or a rival colleague might criticize the bidding team for having missed an opportunity, and the bidding team may want to preemptively defend against such a critique. By selecting a minimax-loss bid the bidding team can claim, "Your alternative bid would have been worse than our bid had there been this other bid distribution. This bid distribution was a real possibility." The minimax bid is then robust to complaints that appeal to the materialized bid distribution.[18] Minimax bids are a way to justify the choice as an (undisputed) counterfactual case can be presented so that the minimax bid was the compromise between the two cases.[19]

Our model can also be understood as featuring (residual) supply uncertainty: let Q be the upper bound of the support of supply and reduce supply through other bidders that demand units at prohibitively high prices, above $v^i(0)$.

[18]Savage [1951] also suggests group decision-making as a justification for the minimax principle. In his story group members have different subjective probability assessments and the minimax principle seeks to keep the greatest "violence" done to anyone's opinion to a minimum. In contrast, we interpret the minimax as a way to defend against ex post complaints.

[19]The worst-case utility is always zero if the bidder seeks to maximize the payoff guarantee. Any bid

3 Two-unit demand

Before our general analysis, we illustrate our analytical approach in the case in which bidders demand up to two units.[20] As in the model of Ausubel et al. [2014], bidder i has value v_{i1} for their first unit and value $v_{i2} = \rho v_{i1}$ for their second unit, and submits two bids (b_{i1}, b_{i2}). We assume marginal values are decreasing and non-negative, i.e., $v_{i1} \geq v_{i2} \geq 0$.[21] Lemma 1 in Section 4 reduces the interim loss-minimization problem to an ex post regret-minimization problem; therefore, it suffices to study opponent bids instead of richer bid distributions. There are three relevant outcomes in an auction in which a bidder demands up to two units: the bidder may win zero, one, or two units. And, if a bidder wins $k \in \{0, 1, 2\}$ units, the opponent's bids can be such that it would have been ex post optimal to win $k' \in \{0, 1, 2\}$ units. Thus, for each k we find the opponent bids (and k') such that bidder i leaves the most surplus on the table—that is, the opponent bids that maximize bidder i's regret conditional on winning k units.

3.1 Pay-as-bid auctions

With decreasing marginal values, it is never profitable to bid above value in a pay-as-bid auction. Hence, we restrict attention to $b_{i1} \leq v_{i1}$ and $b_{i2} \leq v_{i2}$.

Case 1: zero units. The two highest opponent bids, denoted by c_1 and c_2 (where $c_1 \geq c_2$), must be above bidder i's bids when bidder i does not win anything: $c_2 \geq b_{i1}$. Let $(x)_+ = \max\{0, x\}$. The highest possible expected utility given c_1 and c_2 is

$$\max\{(v_{i1} - c_2)_+, (v_{i1} - c_1)_+ + (v_{i2} - c_1)_+\};$$

it can be optimal to win one unit by bidding marginally higher than c_2 (the left-hand term), or it can be optimal to win two units by having both bids marginally higher than c_1 (the right-hand term). However, if $c_2 \geq v_{i1}$, then it is optimal to not win anything and the highest

function below value is then optimal. A natural selection is to bid value (because it is the only bid function for which no opponent bid would induce a change in the bid function, namely to overbid the opponent slightly). Bidding value is neither a good predictor of behavior (*cf.* the experimental data of Engelmann and Grimm [2009]), nor a sensible bid recommendation because there is no point in bidding (at least in the discriminatory auction). Put differently, maxmin expected utility is not robust to complaints about missed opportunities.

[20] An earlier working version of this paper contains the analysis of the general discrete multi-unit case [Kasberger and Woodward, 2023].

[21] We analyze increasing marginal values in Appendix B.

possible expected utility is zero. Maximal regret conditional on winning nothing is therefore

$$\sup_{(c_1,c_2):\, c_1 \geq c_2 > b_{i1}} \underbrace{\left[\max\left\{(v_{i1}-c_2)_+, (v_{i1}-c_1)_+ + (v_{i2}-c_1)_+\right\}\right]}_{\text{highest possible surplus given opponent bids}} - \underbrace{[0]}_{\text{surplus with original bids}}.$$

Observe that regret cannot be maximized at $c_2 \geq v_{i1}$ because this would lead to zero regret, implying that the original bids were optimal. Then regret is decreasing in at least one of c_1 and c_2, and in the worst case (for bidder i) the opponent's bids are $c_1 = c_2 = b_{i1} + \varepsilon$. Depending on b_{i1} and v_{i2}, the bidder wants to win one or two units and can win those by raising the first bid to $b_{i1} + 2\varepsilon$ and (if profitable) the second bid to $b_{i1} + 2\varepsilon$. Put differently, bids are most suboptimal if the bidder could have won as many as they wanted at a price just above their first-unit bid. In the limit ($\varepsilon \to 0$), worst-case regret conditional on winning zero units equals

$$(v_{i1} - b_{i1}) + (v_{i2} - b_{i1})_+ .$$

Case 2: one unit. The bidder wins one unit if $b_{i1} > c_2$ and $c_1 > b_{i2}$. Conditional on this outcome, maximal regret is

$$\sup_{(c_1,c_2):\, b_{i1} > c_2 \text{ and } c_1 > b_{i2}} \underbrace{\left[\max\left\{v_{i1}-c_2, (v_{i1}-c_1)_+ + (v_{i2}-c_1)_+\right\}\right]}_{\text{highest possible surplus given opponent bids}} - \underbrace{[v_{i1} - b_{i1}]}_{\text{surplus with original bids}}.$$

The highest possible surplus may be attained by winning one or two units (winning nothing cannot be optimal since bids are below value). When it is optimal to win a single good, regret is maximized when this can be done at a lower price. In the worst case, $c_1 \geq v_{i2}$ and $c_2 = 0$, and associated regret is b_1—i.e., the bidder overpays for the unit they receive. But note that regret in this case is no higher than in the case in which the bidder receives two units, and ex post prefers to receive two units.

Then it is sufficient to consider only the possibility that, given c_1 and c_2, bidder i still desires to win two units but at a different price. In this case, regret decreases in c_1. In the worst case the opponents' highest bid is $c_1 = b_{i2} + \varepsilon$. The bidder overpays on the first unit because a lower first bid would have also been winning, and bids marginally too low on the second unit. We refer to this case as *underbidding* because the bidder bids too little on the second unit. Worst-case regret conditional on winning one unit is

$$[(v_{i1} - b_{i2}) + (v_{i2} - b_{i2})] - [(v_{i1} - b_{i1})] = (b_{i1} - b_{i2}) + (v_{i2} - b_{i2}) .$$

Case 3: two units. The bidder wins two units if $b_{i2} > c_1$. Conditional on this outcome, maximal regret is

$$\sup_{(c_1,c_2):\, b_{i2}>c_1} [\max\{(v_{i1}-c_1)+(v_{i2}-c_1), v_{i1}-c_2\}] - [(v_{i1}-b_{i1})+(v_{i2}-b_{i2})].$$

Conditional on winning both units, bids can be suboptimal if they are too high; the bidder overpays on both units if both units could be acquired cheaper (for a per-unit price of $c_1+\varepsilon$) and overpays on the first unit if one unit can be acquired cheaper (for a per-unit price of $c_2+\varepsilon$); as when the bidder received one unit, when bids are below values it cannot be the case that it is ex post optimal to receive zero units. In the worst case, the two highest opponent bids are $(0,0)$ so that the bidder can reduce their bids to $(\varepsilon,\varepsilon)$; in this event worst-case regret equals

$$[(v_{i1}-0)+(v_{i2}-0)] - [(v_{i1}-b_{i1})+(v_{i2}-b_{i2})] = b_{i1}+b_{i2}.$$

Determining optimal bids. The minimax-loss bid balances the regret conditional on the realization of any of the three outcomes: underbidding regret conditional on receiving any number of units, and overbidding regret conditional on winning two units. Maximal loss is

$$\max\left\{(v_{i1}-b_{i1})+(v_{i2}-b_{i1})_+\,,\ (b_{i1}-b_{i2})+(v_{i2}-b_{i2})\,,\ b_{i1}+b_{i2}\right\}. \tag{1}$$

The minimax is attained by equalizing the three expressions, and the unique minimax-loss bid vector in the pay-as-bid auction is

$$b_{i1}^{\text{PAB}} = \begin{cases} \frac{v_{i1}}{9}(3+2\rho) & \text{if } \rho \geq \frac{3}{7}, \\ \frac{v_{i1}}{6}(3-\rho) & \text{if } \rho < \frac{3}{7}; \end{cases} \quad \text{and} \quad b_{i2}^{\text{PAB}} = \frac{\rho v_{i1}}{3}. \tag{2}$$

The case distinction is due to the value for the second good being below or above the bid for the first; i.e., the term $(v_{i2}-b_{i1})_+$ in Equation (1).

Figure 1 illustrates the bid function as a function of ρ. If $\rho=0$, then the minimax bid is $b_{i1}^{\text{PAB}} = v_{i1}/2$, which is as in the first-price auction for a single good [Kasberger and Schlag, 2024]. The bid b_{i1}^{PAB} *decreases* in ρ for $\rho \leq 3/7$.[22] This antitonicity arises because increasing ρ in this range increases the loss conditional on receiving a single unit, hence the bid b_{i1} falls so that loss is equalized across outcomes. To provide more discussion of this antitonicity, suppose ρ is relatively low so that the bidder does not want to buy the second good at price b_{i1}. The bid b_{i1} is then found by equalizing the underbidding regret conditional on losing the

[22] McAdams [2007] provides examples of a uniform-price auction where Bayes-Nash equilibrium bids may decrease in the bidder's value due to risk aversion and affiliated values. Our example concerns the pay-as-bid auction and relies on a distinct rationale.

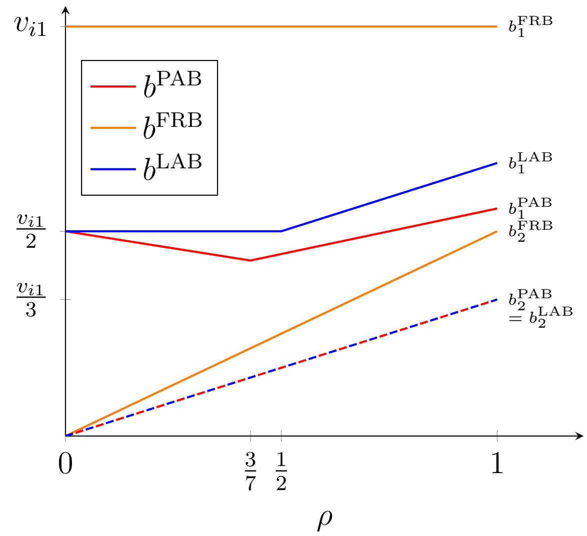

Figure 1: First- and second-unit bids in the pay-as-bid and uniform-price auctions, when the bidder demands two units.

auction with the overbidding regret conditional on winning two units: $v_{i1} - b_{i1} = b_{i1} + b_{i2}$. A higher b_{i2} increases the level of regret, forcing b_{i1} to be lower. Since $b_{i2} = v_{i2}/3$ increases in ρ, the observed antitonicity follows. Conversely, if ρ is sufficiently high so that the bidder wants to buy two units at a price of b_{i1}, underbidding regret conditional on losing the auction is $(1 + \rho)v_{i1} - 2b_{i1}$. A marginal increase of ρ now increases regret conditional on losing the auction stronger than overbidding regret. Hence, to maintain equivalence between overbidding and underbidding regret, the bid b_{i1} must increase in ρ. It follows that for values of ρ above $3/7$, both bids increase in ρ, though the second bid b_{i2}^{PAB} increases more quickly than b_{i1}^{PAB}. By corollary, the spread between the two bids uniformly decreases in ρ.

3.2 Uniform-price auctions

As in the pay-as-bid auction, three exhaustive outcomes may maximize loss in the uniform-price auction: the bidder either receives zero, one, or two units. We consider these outcomes on a case-by-case basis and study the uniform-price auction with the first rejected bid (FRB) or the last accepted bid (LAB) as the market-clearing price.

3.2.1 First rejected bid uniform-price auction

In the uniform-price auction with the first rejected bid as the market-clearing price, it is known that bidding truthfully on the first unit is weakly dominant [Engelbrecht-Wiggans and Kahn, 1998]. However, to motivate a selection of a minimax-loss bid we compute maximal regret for any undominated bids, $b_{ij} \leq v_{ij}$, $j = 1, 2$.

Case 1: zero units. Recall that bidder i wins nothing if the second-highest competing bid is above b_{i1}: $c_2 \geq b_{i1}$. The difference from the pay-as-bid auction lies in the payment rule. If bidder i were to win one unit, then the first rejected bid would be $\max\{c_2, b_{i2}\}$. Hence, the maximum ex post utility in which bidder i wins one unit is associated with setting the second bid equal to zero. If bidder i were to win two units, then the first rejected bid would be c_1. Maximal regret conditional on winning zero units is

$$\sup_{(c_1,c_2):\, c_1 \geq c_2 > b_{i1}} [\max\{(v_{i1} - c_1)_+ + (v_{i2} - c_1)_+, (v_{i1} - c_2)_+\}] - [0].$$

Regret decreases in c_1 and c_2 and is maximized by $c_1 = c_2 = b_{i1} + \varepsilon$. Maximal regret is therefore

$$\max\left\{(v_{i1} - b_{i1}) + (v_{i2} - b_{i1})_+, (v_{i1} - b_{i1})_+\right\}.$$

This case is associated with "underbidding" because bidder i could win more by raising the first bid marginally, and the second bid if profitable.

Case 2: one unit. Recall that the bidder wins one unit if $b_{i1} > c_2$ and $c_1 > b_{i2}$. The first rejected bid is $\max\{c_2, b_{i2}\}$. Conditional on this outcome, maximal regret is

$$\sup_{(c_1,c_2):\, b_{i1} > c_2 \text{ and } c_1 > b_{i2}} [\max\{(v_{i1} - c_2)_+, (v_{i1} - c_1)_+ + (v_{i2} - c_1)_+\}] - [v_{i1} - \max\{c_2, b_{i2}\}].$$

Given c_1 and c_2, it can be optimal to win one or two units, but it cannot be optimal to win zero because bids are below value. When it is optimal to win one unit, regret equals

$$\sup_{(c_1,c_2): b_{i1} > c_2 \text{ and } c_1 > b_{i2}} v_{i1} - c_2 - v_{i1} + \max\{c_2, b_{i2}\} = b_{i2}.$$

Corresponding worst-case opponent bids are $c_1 = \max\{v_{i1}, b_{i1}\}$ and $c_2 = 0$. Bidder i's second bid b_{i2} sets the market-clearing price so that bidder i "overbids" because they could have won the unit for free.

When winning two units is optimal, regret equals

$$\sup_{(c_1,c_2):\, b_{i1} > c_2 \text{ and } c_1 > b_{i2}} v_{i1} - c_1 + v_{i2} - c_1 - [v_{i1} - \max\{c_2, b_{i2}\}].$$

If $c_2 > b_{i2}$, then regret is $v_{i2} + c_2 - 2c_1$. Since regret increases in c_2 and $c_1 \geq c_2$, it follows that maximal regret in this case is $\sup_{c_1 > b_{i2}} v_{i2} - c_1$. Worst-case regret is $v_{i2} - b_{i2}$; the corresponding worst-case bids are $c_1 = c_2 = b_{i2} + \varepsilon$. If $c_2 \leq b_{i2}$, then regret is $v_{i2} - 2c_1 + b_{i2}$. Regret decreases in c_1 so that the worst-case c_1 equals $b_{i2} + \varepsilon$. In any case, if it is optimal to win two units, the regret is associated with "underbidding;" a marginally higher b_{i2} would have been better. Maximal regret conditional on winning one unit is $\max\{b_{i2}, v_{i2} - b_{i2}\}$.

Case 3: two units. Recall that two units are won if $b_{i2} > c_1$. In this case, the highest rejected bid is always c_1. Hence, the two units cannot be won cheaper. Winning nothing cannot be better than two units if bids are below value. If it is optimal to win one unit for a transfer of c_2, then regret is

$$\sup_{(c_1,c_2):\, b_{i2}>c_1} [v_{i1} - c_2] - [v_{i1} + v_{i2} - 2c_1] = \sup_{(c_1,c_2):\, b_{i2}>c_1} 2c_1 - c_2 - v_{i2}.$$

In the worst case, $c_2 = 0$ and $c_1 = b_{i2} - \varepsilon$ so that maximal regret is $2b_{i2} - v_{i2}$. Bidder i "overbids" because one unit could have been won for free.

Determining optimal bids. Regret is maximized by one of the previous cases and equal to

$$\max\{(v_{i1} - b_{i1}) + (v_{i2} - b_{i1})_+, b_{i2}, v_{i2} - b_{i2}, 2b_{i2} - v_{i2}\}. \tag{3}$$

Regret minimization pins down $b_{i2} = v_{i2}/2$. However, in contrast with the pay-as-bid auction, there is no case that involves both b_{i1} and b_{i2}. Consequently, there is no unique bid that minimizes the maximal regret.[23] A natural minimax-bid is found "pointwise" by selecting bids that minimize the maximum of the expressions in which they appear. A truthful first bid $b_{i1} = v_{i1}$ is then optimal (as are marginally lower first bids). The minimax-loss bid is

$$b_{i1}^{\text{FRB}} = v_{i1} \quad \text{and} \quad b_{i2}^{\text{FRB}} = \frac{\rho v_{i1}}{2}. \tag{4}$$

Below we call this selection the *conditional regret minimizing bid.*

We view this selection as natural for three reasons. First, it leads to a bidding strategy that is weakly dominant [Engelbrecht-Wiggans and Kahn, 1998]. Second, the selection is unique, serving the purpose of pinning down the bidding strategy. Third, it optimizes the entire bidding function also for "local worst cases". To elaborate, let $b_{i1} = v_{i1} - \varepsilon$ and $b_{i2} = v_{i2}/2$. Such a bid also minimizes maximal regret, which is equal to $v_{i2}/2$. A "global worst case" for such a bidding function is that bidder i wins one unit and pays b_{i2}, but could have won the unit for free. However, the bidding function does not adequately protect against the situation in which bidder i wins zero units but could have won at least one with a higher bid on the first unit. The conditional regret is relatively minor (e.g., $v_{i1} - b_{i1} = \varepsilon$) compared to the global maximal regret of $v_{i2}/2$. Nevertheless, it might pay off to minimize maximal loss also for "local worst cases" instead of only the global ones. Note that such a

[23] The best-reply correspondence of Bayesian bidders can also be multi-valued. Moreover, multiple equilibria can exist. For example, in the "collusive" equilibrium with two bidders and units, bidder i bids truthfully on the first unit and zero on the second. An alternative equilibrium is that bidder i bids some high value $\overline{v}$ for the first unit and zero for the second unit.

bid would not necessarily be separable with respect to v_{i1} and v_{i2}.

3.2.2 Last accepted bid uniform-price auction

There is also not a unique minimax-loss bid in the last accepted bid uniform-price auction. The conditional regret minimizing bid equals

$$b_{i1}^{\text{LAB}} = \begin{cases} \frac{1}{3} v_{i1} \left(1+\rho\right) & \text{if } \frac{1}{2} \leq \rho, \\ \frac{1}{2} v_{i1} & \text{otherwise;} \end{cases} \quad \text{and} \quad b_{i2}^{\text{LAB}} = \frac{\rho v_{i1}}{3}. \tag{5}$$

We provide the details in Appendix D.

3.3 Revenue and welfare comparison

We now use the minimax-loss bid functions to compare auction revenue and welfare across the three auction formats. Despite the clear ranking of bid levels (*cf.* Figure 1), auction outcomes may be ambiguous because different payment rules imply different bid functions which induce distinct mappings from bidder values to outcomes. Minimax-loss bids are insensitive to the underlying distribution of bidder values, and this distribution therefore induces a degree of freedom which may render cross-mechanism outcome comparisons ambiguous. Essentially, the uniform-price formats generate higher revenue when the distribution of values is narrow and pay-as-bid generates higher revenue when the distribution of values is wide; this is because in the pay-as-bid auction the winning bidders' payments are independent of opponent competitiveness, while in the FRB uniform-price auction the lack of competition affects the price paid.

We illustrate these tradeoffs in the two-bidder, two-unit context of this section.

Example 1. Let there be two bidders and two units for sale.[24] The marginal values are flat (i.e., $\rho = 1$ and $v_{i2} = v_{i1} \equiv v_i$). Without loss of generality, assume that bidder 1's marginal value is higher than bidder 2's marginal value, $v_1 \geq v_2$, and let $\tau \in [0,1]$ be such that $v_2 = \tau v_1$.

In the PAB auction, bidder i bids $(\frac{5}{9} v_i, \frac{1}{3} v_i)$. Bidder 1 wins two units if $\frac{1}{3} v_1 \geq \frac{5}{9} \tau v_1$, that is, if and only if $\tau \leq \frac{3}{5}$. Otherwise, both bidders win one unit each. Ex-post revenue in the PAB auction is $\frac{8}{9} v_1$ if $\tau \leq \frac{3}{5}$ and $\frac{5}{9} v_1 (1+\tau)$ if $\tau > \frac{3}{5}$.

In the FRB uniform-price auction, bidder i bids $(v_i, \frac{1}{2} v_i)$. Bidder 1 wins two units if $\frac{1}{2} v_1 \geq \tau v_1$, that is, if and only if $\tau \leq \frac{1}{2}$. In this case, the highest rejected bid is τv_1.

[24] The analysis of the PAB and LAB auctions also applies when there are more than two bidders. However, the first rejected bid might come from the bidder with the third highest bid.

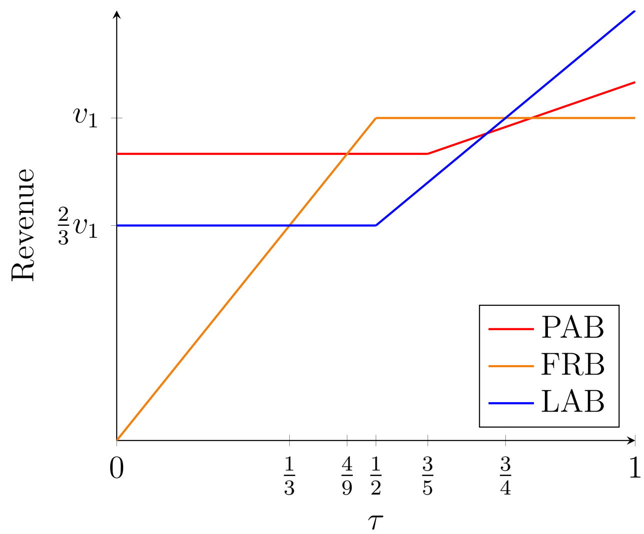

Note: Two bidders and two discrete units for sale. Marginal values are flat. The highest marginal value is denoted by v_1; the second-highest marginal value is $v_2 = \tau v_1$.

Figure 2: Ex-post revenue in Example 1

Otherwise, both bidders win one unit each and the first rejected bid is $v_1/2$. Ex-post revenue in the FRB uniform-price auction is $2\tau v_1$ if $\tau \leq \frac{1}{2}$ and v_1 if $\tau > \frac{1}{2}$.

In the LAB uniform-price auction, bidder i bids $(\frac{2}{3}v_i, \frac{1}{3}v_i)$. Bidder 1 wins two units if $\frac{1}{3}v_1 \geq \frac{2}{3}\tau v_1$, that is, if and only if $\tau \leq \frac{1}{2}$. In this case, the last accepted bid is $\frac{1}{3}v_1$. Otherwise, both bidders win one unit each so that the last accepted bid equals $\frac{2}{3}\tau v_1$. Ex-post revenue in the LAB uniform-price auction is $\frac{2}{3}v_1$ if $\tau \leq \frac{1}{2}$ and $\frac{4}{3}\tau v_1$ if $\tau > \frac{1}{2}$.

Figure 2 depicts ex-post revenue in the three auction formats as a function of τ. The first observation is that revenue can be highest in any auction format. The PAB auction leads to the highest revenue if the second-highest marginal value is relatively low ($\tau \leq \frac{4}{9}$). Bidder 1 wins both objects and revenue is high due to the relatively high first bid. The FRB uniform-price auction maximizes revenue among the three auction formats if the second-highest marginal value takes intermediate values ($\frac{4}{9} \leq \tau \leq \frac{3}{4}$). Revenue is (close to) v_1, which is higher than (revenue in the neighborhood of) $8v_1/9$ in the PAB and $2v_1/3$ in the LAB uniform-price auction. Finally, the LAB uniform-price auction generates the highest revenue if both bidders have similar values since the last accepted bid is relatively high in this case ($\tau \geq \frac{3}{4}$). The expected revenue of the three auctions can also be ranked analogously according if the bidders' joint value distribution puts sufficient mass on τ in the respective intervals. ◀

Any auction can also lead to the highest welfare. In the notation of Example 1, it is efficient that the bidder with the highest type (v_1) wins two units if $v_1 + \rho v_1 \geq v_1 + \tau v_1$, i.e., if and only if $\rho \geq \tau$. Otherwise, in the efficient allocation the two bidders with the two highest values each receive one unit. If $\rho = 1$, it is always efficient that the bidder with the highest value wins both units. The PAB auction is efficient if and only if $\tau \leq \frac{3}{5}$, while the uniform-price auctions are efficient if and only if $\tau \leq \frac{1}{2}$. It follows that all three auctions are efficient if the distribution of τ puts all the mass below $\frac{1}{2}$ and equally inefficient if all mass is above $\frac{3}{5}$. However, if the distribution of τ puts mass on $(\frac{1}{2}, \frac{3}{5}]$, then PAB is more efficient than the uniform-price auctions. We note that welfare equivalence in the one-unit-to-each case does not contradict the ambiguous revenue ranking discussed above, as welfare measurement only considers the extensive margin (whether the "right" agent receives the good) while revenue measurement also considers the intensive margin (how much they pay).

The following proposition shows that the pay-as-bid auction achieves a weakly higher welfare than the other auctions.

Proposition 1. *Let there be two discrete units for sale and $v_{i2} = \rho v_{i1}$ for all bidders i, $0 \leq \rho \leq 1$. Suppose bidders play the minimax-loss bids in Equations* (2)*,* (4)*, and* (5)*.*

- *All three auction formats are efficient when it is efficient that the two bidders with the highest values win one unit each.*
- *When it is efficient that the bidder with the highest value wins both units, the pay-as-bid auction is weakly more efficient than the last accepted bid uniform-price auction, which is weakly more efficient than the first rejected bid uniform price auction.*

3.4 Comparison to Bayes-Nash equilibrium and design implications

We now compare minimax-loss and Bayes-Nash equilibrium bids and outcomes. We first argue that minimax-loss and BNE bids cannot be generally compared; whether one bid curve or the other is more aggressive depends on the steepness of true demand and on the level of competition in the auction. An immediate corollary is that auction outcomes cannot be compared unambiguously. We therefore cannot say conclusively that an auctioneer should encourage minimax-loss bidding or BNE bidding, to the extent such encouragement is possible. Nonetheless, in spite of this lack of comparability we show that minimax-loss bids are qualitatively consistent with experimental work, hence auctioneers may have reason to seriously consider the implications of minimax-loss bidding.

The comparison of minimax-loss and BNE bid curves depends on the strength of competition and on the strength of demand. Intuitively, when there is little competition in the auction, Bayes-Nash equilibrium bids tend to be lower than bids under maximal uncertainty because only Bayesian bidders adjust their bids to the absence of competition. Maximally uncertain bidders are unaware that competition is thin and, accordingly, tend to bid higher than in BNE. If there is strong competition in the sense of there being many bidders, minimax-loss bids will tend to be lower than Bayes-Nash equilibrium bids because the bid function is independent of the number of bidders while BNE bids tend to value as the number of bidders becomes large [Swinkels, 2001].

To demonstrate the role of demand strength, in what follows we hold the number of bidders equal to two. Even under this minimal level of competition minimax-loss and BNE bids differ qualitatively and cannot be ranked unambiguously. We show this for pay-as-bid auctions under the assumption that values are distributed uniformly. For the case of flat marginal values ($\rho = 1$), Ausubel et al. [2014] show that Bayes-Nash equilibrium bids are flat; that is, each bidder submits two identical bids. This stands in contrast to bidding under maximal uncertainty, where the second bid is strictly lower than the first (Equation (2)). With decreasing marginal values ($0 \leq \rho < 1$), the BNE in the PAB auction is no longer flat [Ausubel et al., 2014] and BNE bids can be higher or lower than those under maximal uncertainty, so the bid functions cannot be ranked uniformly. However, as ρ becomes small, there is effectively no competition as each bidder demands one of two units. Bayesian bidders know this and therefore bid zero in equilibrium. Maximally uncertain bidders do not know that there is no competition and compete more intensely than Bayesian bidders. More formally, both BNE bids converge to zero as $\rho \to 0$, while the first minimax-loss bid converges to the first-price auction minimax-loss bid of $v_{i1}/2$ [Kasberger and Schlag, 2024]. In this setting, the bids under maximal uncertainty are higher than in BNE.

Submitted bids also cannot be unambiguously ranked in uniform-price auctions. In the FRB uniform-price auction, Bayes-Nash equilibria can be collusive in the sense that the second bid is zero; this can even be the unique BNE [Ausubel et al., 2014, Engelbrecht-Wiggans and Kahn, 1998, Noussair, 1995]. In contrast, bidders submit two positive bids under maximal uncertainty because it could be that the other bidder has a very low first bid, so the bidder would regret not submitting a positive second bid. In the LAB uniform-price auction, the minimax-loss and BNE bid sometimes share a property called "separability" and a relation to the bid in the first-price auction. A bid function is separable if the bid for quantity k only depends on the marginal value for the k^{th} unit. Burkett and Woodward [2020a] identify a model in which the BNE bid is separable. Moreover, they show that the equilibrium bid for quantity k is as in a first-price auction with value v_k. In contrast, the minimax-loss bid (5) is separable only if v_{i2} is sufficiently low. In this case, the first bid is then also as in a first-price auction. The second, however, is lower than in a first-price auction with value v_{i2}.

These qualitative findings—that PAB bids are not flat, and that FRB bids are nonzero for the second unit—align with experimental work on multi-unit auctions. Engelmann and Grimm [2009] run laboratory versions of PAB and FRB. Subjects in experimental PAB auctions do not use flat bids but bid spreads that are qualitatively consistent with the minimax-loss approach [Engelmann and Grimm, 2009]. In the FRB experiment of Engelmann and Grimm [2009], subjects rarely bid zero on the second item, which is qualitatively consistent with the minimax-loss approach but not with the equilibrium predictions.

Finally, auction design implications when facing minimax-loss bidders can be in direct contradiction with design implications when facing Bayesian bidders. The existence of zero-revenue Bayes-Nash equilibria suggests that the FRB uniform-price auction might not be a good choice in terms of revenue. However, Figure 2 shows that it leads to the highest revenue among the three auction formats when the second-highest value is about half of the highest value. For example, suppose marginal values are constant and perfectly correlated with $v_2 = v_1/2 + \varepsilon$ for ε positive but small. There clearly exists a zero-revenue equilibrium in the FRB auction,[25] but the FRB auction is revenue-optimal among the three auction formats under maximal uncertainty. In this setting, the FRB auction is also efficient with maximally uncertain bidders, while the zero-revenue equilibrium is inefficient. Moreover, in a setting nested by Proposition 1, Figure 2 of Ausubel et al. [2014] shows that the PAB auction leads to the lowest expected surplus when ρ is low. We find, however, that the PAB auction is weakly welfare-dominant among the three formats.

4 Loss in auctions for homogeneous goods

We now establish general properties of the minimax-loss problem. Under maximal uncertainty, bidders believe every possible distribution of opponent bids is feasible. Since degenerate distributions are believed to be feasible and turn out to maximize loss, maximum loss is equivalent to maximum regret. This is a consequence of the linearity of bidder preferences and not specific to the analysis of auctions or other features of our model.

Lemma 1 (Reduction to maximum regret)**.** *Under maximal uncertainty, maximizing loss is equivalent to maximizing regret. That is, for all values v^i and bids b^i,*

$$\sup_{B^{-i}\in\mathcal{B}} L\left(b^i; B^{-i}, v^i\right) = \sup_{b^{-i}} R\left(b^i; b^{-i}, v^i\right).$$

[25] Bidder 1's utility from winning two units is $2v_1 - 2v_1/2 - 2\varepsilon = v_1 - 2\varepsilon$, while it is v_1 when winning one unit for free.

To simplify the regret maximization problem, we decompose it to the related problem of maximizing *conditional regret*. Given any quantity $q \in [0, Q]$, let $R_q(b^i; v^i)$ denote bidder i's (maximal) regret conditional on winning q units. Given a bid b^i and an opponent bid b^{-i}, bidder i's quantity allocation $q^i(b^i, b^{-i})$ is deterministic. Since maximum loss is identical to maximum regret, which is derived ex post after opponent demand is realized, it follows that maximum loss is the highest conditional regret from receiving any quantity, $\sup_{b^{-i}} R(b^i; b^{-i}, v^i) = \sup_q R_q(b^i; v^i)$. Conditional regret forms the basis of our subsequent analysis.

4.1 Pay-as-bid auctions

To develop intuition for loss maximization in the pay-as-bid auction, consider the potential sources of regret in a canonical single-unit first-price auction. Ex post, bids in single-unit discriminatory auctions are either too high—because the bidder strictly outbids the second-highest bidder—or too low—because the bidder underbids the highest bidder, whose bid was below the bidder's value [Kasberger and Schlag, 2024].[26] This same intuition is true pointwise in multi-unit pay-as-bid auctions: the bidder frequently would prefer to increase their bid for large quantities and decrease their bid for small quantities. We use this observation to pin down conditional regret in the pay-as-bid auction.

If bidder i submits bid b^i and obtains quantity q, they know that the market-clearing price is $p^\star \in [b^i_+(q), b^i(q)]$, where $b^i_+(q) = \lim_{q' \searrow q} b^i(q')$.[27] Their regret is at least

$$\underline{R}_q^{\text{PAB}}\left(b^i; p^\star, v^i\right) = \int_0^q \left(b^i(x) - p^\star\right) dx + \int_q^Q \left(v^i(x) - p^\star\right)_+ dx.$$

That is, their regret is at least their overpayment for units they received, plus the utility foregone by underbidding for units they value above the market-clearing price. This regret would be realized if, for example, all opponents submitted flat bids at the price $p^\star$. This expression is strictly decreasing in $p^\star$; hence, bidder i's conditional regret is at least

$$\underline{R}_q^{\text{PAB}}\left(b^i; v^i\right) = \int_0^q \left(b^i(x) - b^i(q)\right) dx + \int_q^Q \left(v^i(x) - b^i(q)\right)_+ dx, \tag{6}$$

since $b^i_+(q) = b^i(q)$ for right-continuous bid functions. Because $\underline{R}_q^{\text{PAB}}$ is the regret the bidder has in the case in which they wish they had bid slightly more for larger quantities, we refer to $\underline{R}_q^{\text{PAB}}$ as *underbidding regret*. The second term can also be written as $\int_q^{{v^i}^{-1}(b^i(q))} v^i(x) -$

[26] When bids are neither too high nor too low, regret is zero. Generally, maximal regret will be nonzero.
[27] For notational simplicity we define $b^i_+(Q) = 0$.

$b^i(x)dx$, where $v^{i^{-1}}$ is the right inverse of v^i and at most equal to Q:

$$v^{i^{-1}}(y) = \sup\{x \in [0, Q] \colon v^i(x) \geq y\}.$$

Alternatively, bidder i might be able to obtain the same allocation by bidding just above zero for all units. This will be the case when their opponents, in aggregate, submit extremely high bids for $Q-q$ units and zero bids for all remaining units. In this case all nonzero payment is wasted, and regret is

$$\overline{R}_q^{\text{PAB}}\left(b^i; v^i\right) = \int_0^q b^i\left(x\right) dx.$$

Because $\overline{R}_q^{\text{PAB}}$ is the regret the bidder has in the case in which they wish they had bid nearly zero for all units, we refer to $\overline{R}_q^{\text{PAB}}$ as *overbidding regret*.

The conditional regret for quantity q is

$$R_q^{\text{PAB}}(b^i; v^i) = \max\left\{\overline{R}_q^{\text{PAB}}\left(b^i; v^i\right), \underline{R}_q^{\text{PAB}}\left(b^i; v^i\right)\right\}.$$

Since $\underline{R}_Q^{\text{PAB}}(b^i; v^i) = \overline{R}_Q^{\text{PAB}}(b^i; v^i)$ and $\overline{R}_q^{\text{PAB}}(b^i; v^i)$ is weakly increasing in q, maximum loss is the supremum of underbidding regret, taken over all quantities q.

Lemma 2 (Maximum loss in pay-as-bid)**.** *In the pay-as-bid auction, maximal loss given bid b^i is*

$$\sup_q \underline{R}_q^{PAB}\left(b^i; v^i\right).$$

A proof is given in Appendix A.2.

4.2 Uniform-price auctions

We first establish expressions for underbidding and overbidding regret in the uniform-price auction, in line with our analysis of pay-as-bid auctions. The market-clearing price is $p^\star \in \{p^{\text{LAB}}, p^{\text{FRB}}\}$. In spite of the potentially large difference in market prices, the strategic analyses of FRB and LAB differ only in the discrete multi-unit case (Section 3).

In the uniform-price auction, bids above the market-clearing price are relevant only to the extent that they guarantee a unit is awarded; they do not otherwise affect the bidder's utility. This is in contrast to the pay-as-bid auction, where bids above the market-clearing price are paid whenever the unit is awarded. When bidder i receives quantity q, the market-clearing price must be $p^\star = b^i(q)$. Bidder i's underbidding regret is

$$\underline{R}_q^{\text{UPA}}\left(b^i; v^i\right) = \int_q^Q \left(v^i\left(x\right) - b^i\left(q\right)\right)_+ dx.$$

As in the pay-as-bid auction, underbidding regret accounts not only for the fact that the bidder might regret not bidding just above the market-clearing price, but also for the fact that the bidder might affect their own transfer.

Alternatively, bidder i might be able to obtain the same allocation by bidding just above zero for all units. This will be the case when their opponents submit high bids for $Q - q$ units and submit zero bids for all remaining units. In this case all nonzero bids are wasted, and regret is higher the higher is the market-clearing price, hence overbidding regret is

$$\overline{R}_q^{\text{UPA}} \left(b^i; v^i\right) = q b^i \left(q\right).$$

This differs from overbidding regret in the pay-as-bid auction, $\overline{R}_q^{\text{PAB}}$, since in the uniform-price auction only the marginal bid is relevant.

The conditional regret for any quantity q is

$$R_q^{\text{UPA}} \left(b^i; v^i\right) = \max \left\{ \underline{R}_q^{\text{UPA}} \left(b^i; v^i\right), \overline{R}_q^{\text{UPA}} \left(b^i; v^i\right) \right\}.$$

Because maximum loss is equal to maximum regret, and ex post regret is obtained at some allocation, maximal loss may be identified with maximizing conditional regret.

Lemma 3 (Maximum loss in uniform-price)**.** *In the uniform-price auction, maximal loss given bid b^i is*

$$\sup_q R_q^{UPA} \left(b^i; v^i\right).$$

5 Unconstrained minimax-loss bids

We now characterize minimax-loss bids when the auctioned good is perfectly divisible. In contrast with the next section, the bidders are not constrained in their number of bids; they can submit any weakly decreasing, weakly positive, continuous bid functions. At the end of the section, we compare the minimax-loss bid functions to those that form a BNE and discuss design implications.

5.1 Pay-as-bid auctions

Recall that Lemma 2 states that loss is maximized by maximizing underbidding regret. To minimize the highest underbidding regret across all quantities q, observe that underbidding regret for quantity q increases in the bids for quantities $q' < q$, decreases in the bid for quantity q, and is unaffected by the bids for quantities $q'' > q$. It follows that if b^i is an optimal bid function, then underbidding regret must be constant at all quantities q.

Lemma 4 (Equal conditional regret in pay-as-bid)**.** *If b^i is a minimax-loss bid function in the pay-as-bid auction, then $\underline{R}_q^{PAB}(b^i; v^i) = \underline{R}_{q'}^{PAB}(b^i; v^i)$ for all $q, q' \in [0, Q]$.*

Lemma 4 gives a straightforward method for computing minimax-loss bids: minimize conditional regret for any quantity, subject to equal conditional regret across all quantities. Although computationally straightforward, optimal bids do not admit a general analytical form. The formula for conditional regret contains an integral over all units which are valued more than a given bid, but the range of integration depends not only on the bidder's values but also on the prospective bid, which complicates the relationship between bid and loss.

The equal conditional regret condition requires the derivative of conditional regret to equal zero, which leads to the differential equation in the following theorem. Regarding the boundary condition, regret conditional on receiving the maximum possible allocation is $\int_0^Q b^i(x)dx$. The fundamental theorem of differential equations implies that solutions to the system cannot cross, hence the bid for quantity Q must be minimal.

Theorem 1 (Unconstrained pay-as-bid bids)**.** *The unique minimax-loss bid in the unconstrained pay-as-bid auction, b^{PAB}, solves*

$$v^i(q) - b^{PAB}(q) = -{v^i}^{-1}\left(b^{PAB}(q)\right)\frac{db^{PAB}}{dq}(q), \text{ s.t. } b^{PAB}(Q) = 0. \tag{7}$$

The minimax-loss bid is strictly below marginal values and strictly decreasing in quantity q.

If $v^i(Q) \geq \frac{1}{Q}\int_0^Q v^i(x)e^{-\frac{x}{Q}}dx$, then differential equation (7) can be solved analytically and the minimax-loss bid in the unconstrained divisible-good pay-as-bid auction equals

$$b^{PAB}(q) = \frac{1}{Q}e^{\frac{q}{Q}}\int_q^Q v^i(x)e^{-\frac{x}{Q}}dx. \tag{8}$$

The differential equation defining minimax-loss bids in the pay-as-bid auction is similar to the first-order condition defining best responses in a standard Bayes-Nash equilibrium; see, e.g., Hortacsu and McAdams [2010], Woodward [2021], and Pycia and Woodward [2025]. The distinction is that in Bayes-Nash equilibrium the first-order condition contains probabilistic effects—increasing the bid for a particular quantity increases the probability that this quantity is received—while the differential equation in Theorem 1 does not. Intuitively, this is because regret is an ex post concept.

The minimax-loss bid is unique for any marginal value function. Uniqueness simplifies the estimation of bidders' private values if one believes that observed (continuous) bid data is generated by bidders playing the minimax-loss bids under maximal uncertainty. In this case, one can infer bidder i's values from the first-order condition in Equation (7). In contrast to the approach relying on BNE equilibrium as the data-generating model, estimating values from minimax-loss bids does not require the difficult estimation of the (opponent) bid distribution [Hortacsu and McAdams, 2010]. From a normative perspective, a unique bid is attractive as it saves one from further assessing the relative merits of all minimax-loss bids.

The tractability of the minimax-loss bid stands in stark contrast to the typical intractability of the BNE in the pay-as-bid auction. As discussed in the introduction, BNE characterizations exist only in relatively simple (usually complete information or one-parameter) economic settings. Theorem 1 provides a closed-form solution for the minimax-loss bidding function for sufficiently flat marginal values. The following example illustrates this case. On the other hand, it is relatively straightforward to numerically compute minimax-loss bids.

Example 2. Let $v(q) = \theta - \rho \cdot q$, where $0 \leq \rho \leq \theta/(2Q)$; the bidder has flat marginal values if $\rho = 0$ and (non-satiated) quadratic utility if $\rho > 0$. The bid function in Equation (8) equals

$$b^{\text{PAB}}(q) = \underbrace{\theta - \rho q}_{=v(q)} - \rho Q - (\theta - 2\rho Q) e^{\frac{q}{Q} - 1}.$$

The upper bound $\rho \leq \theta/(2Q)$ guarantees that marginal values are sufficiently flat: $v(Q) \geq b^{\text{PAB}}(0)$. In this case, $v^{-1}(b^{\text{PAB}}(q)) = Q$ for any q and Equation (7) can be solved analytically. The bidding function reveals that bidders "shade" their marginal values. The bidding function is non-linear in q unless $\theta = 2\rho Q$. In the latter case, $b^{\text{PAB}}(q) = \frac{\theta}{2Q}(Q - q)$. ◀

5.2 Uniform-price auctions

We now analyze the unconstrained uniform-price auction. Following Lemma 3, maximum loss is

$$L^{\text{UPA}}\left(b^i; v^i\right) = \sup_q R_q^{\text{UPA}}\left(b^i; v^i\right) = \sup_q \max\left\{\underline{R}_q^{\text{UPA}}\left(b^i; v^i\right), \overline{R}_q^{\text{UPA}}\left(b^i; v^i\right)\right\}.$$

That is, maximum loss is a maximum over conditional regrets, which are defined as the higher of overbidding and underbidding regrets for quantity q. Importantly, the bid for quantity q only appears in the conditional regrets for quantity q: $\underline{R}_q^{\text{UPA}}(b^i; v^i)$ and $\overline{R}_q^{\text{UPA}}(b^i; v^i)$. Since $\underline{R}_q^{\text{UPA}}$ is decreasing in $b^i(q)$ and $\overline{R}_q^{\text{UPA}}$ is increasing in $b^i(q)$, if b^i is a minimax-loss bid function, then there must be some quantity q so that $\underline{R}_q^{\text{UPA}}(b^i; v^i) = \overline{R}_q^{\text{UPA}}(b^i; v^i)$. Moreover, regret minimization pins down the bids only for quantities for which conditional regret is maximal. For other quantities, conditional regret minimization leaves the bids (partially) indeterminate. The following proposition is immediate.

Proposition 2 (No unique optimal bid in uniform-price)**.** *There is not a unique minimax-loss bid in the unconstrained uniform-price auction.*

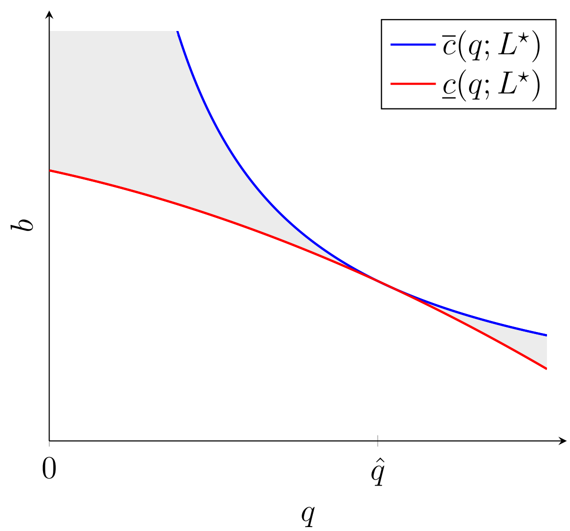

Figure 3: Iso-loss curves of unconstrained underbidding and overbidding regret in the uniform-price auction.

The partial indeterminacy can be illustrated by two *iso-loss curves*. Given loss L, the upper iso-loss curve is $\overline{c}(\cdot; L)$ such that $q\overline{c}(q; L) = L$, and the lower iso-loss curve is $\underline{c}(\cdot; L)$ such that $\int_q^Q (v^i(x) - \underline{c}(q; L))_+ dx = L$. The bid $b(q) = \overline{c}(q; L)$ induces overbidding loss which is constant in quantity, and the bid $b(q) = \underline{c}(q; L)$ induces underbidding loss which is constant in quantity.[28] Figure 3 illustrates the two iso-loss curves. The upper iso-loss curve is always a hyperbola; the lower iso-loss curve depends on marginal values. Bids above the upper iso-loss curve induce loss above L by inducing overbidding regret above L, and bids below the lower iso-loss curve induce loss above L by inducing underbidding regret above L. It follows that the minimax-loss bid must lie entirely between the upper and lower iso-loss curves.

Figure 3 illustrates the upper and lower iso-loss curves for a loss-level equal to minimax loss. In the unconstrained case the upper and lower iso-loss curves are tangent to each other. The bids at the points of tangency are uniquely determined and equal to the conditional regret minimizing bids. In the example depicted in the figure, there is a single point of tangency $\hat{q}$. Other bids are only partially determined; any bid must be below the upper iso-loss curve and above the lower iso-loss curve. In the figure any decreasing bidding function in the shaded area is a minimax bid. All minimax bidding functions agree at $\hat{q}$.

The multiplicity of minimax-loss bids vanishes when strengthening the decision criterion similar to requiring subgame perfection in Nash equilibria of extensive form games. Recall that the multiplicity arises because regret is globally maximized at a single quantity $\hat{q}$ (as in

[28]The same logic does not apply to the pay-as-bid auction, since overbidding regret is monotonically increasing in quantity.

Figure 3). Hence, if player i plays any minimax-loss strategy, any worst-case bid distribution is such that bidder i wins $\hat{q}$. Strengthening the decision criterion to off-path robustness then requires that even if non-worst-case quantity $q \neq \hat{q}$ is won, regret conditional on winning q is minimized. More generally, the perfection requirement is that even in outcomes that do not arise in the worst case, the player minimizes conditional regret. In our case, a *conditional regret minimizing* strategy minimizes the larger of overbidding regret for quantity q and underbidding regret for quantity q. Since regret is maximized by conditional regret for some quantity, a conditional regret-minimaxing bid is a minimax-loss bid and hence a selection of the minimax-loss correspondence.

Definition 1. *The minimax-loss bid b^i is a* conditional regret minimizing *bid if $\underline{R}_q^{UPA}(b^i; v^i) = \overline{R}_q^{UPA}(b^i; v^i)$ for all $q \in [0, Q]$.*

The appeal of conditional regret minimizing bids is that any bid $b^i(q)$ is justifiable ex post. If another minimax bid was chosen so that the bid for quantity q was below the respective conditional regret minimizing bid for that unit, then after winning q units, the case can be made that this bid was too low as it would have been profitable to win more units. Only the conditional regret minimizing bid does not allow such complaints as the regret of paying too much for q units serves as a defense.

Conditional regret minimization requires

$$qb^{\text{UPA}}(q) = \int_q^Q \left(v^i(x) - b^{\text{UPA}}(q)\right)_+ dx$$

for all $q \in [0, Q]$. The conditional regret minimizing bid is unique because overbidding regret increases in bid while underbidding regret decreases in bid.

Theorem 2 (Conditional regret minimizing bid in unconstrained uniform-price auction)**.** *In the unconstrained uniform-price auction, there is a unique conditional regret minimizing bid, b^{UPA}, and this bid solves*

$$qb^{UPA}(q) = \int_q^Q \left(v^i(x) - b^{UPA}(q)\right)_+ dx, \ \forall q \in [0, Q].$$

The theorem implies that $b^{\text{UPA}}(0) = v^i(0)$, i.e., it is optimal to bid value for the "first unit." Moreover, it is optimal to bid 0 for Q: $b^{\text{UPA}}(Q) = 0$.

Although the bidding function of Theorem 2 cannot be compared to all Bayes-Nash equilibria of the uniform-price auction, it is apparent that it does not resemble "collusive" low-revenue equilibria that are frequently discussed in the literature [Ausubel et al., 2014, Marszalec et al., 2020].[29] Indeed, only the bid on the last unit (Q) is zero in a conditional regret minimizing strategy under maximal uncertainty, while many bids are zero in the canonical low-revenue Bayes-Nash equilibrium.

As in the pay-as-bid auction (Example 2), there are convenient expressions for the minimax-loss bids in the uniform-price auction when marginal utility is flat or linear.

Example 3. Let $v(q) = \theta - \rho \cdot q$, where $0 \leq \rho \leq \theta/Q$. The marginal value function is as in Example 2, with the exception that ρ can now take larger values; the new constraint on ρ ensures that the utility function is non-satiated on $[0, Q]$. We distinguish two cases. In the first case, $v^{-1}(b(q))$ is less than Q, which is true for q close to 0 (because these bids are close to $v(0)$). The conditional regret minimizing bid then solves

$$b \cdot q = \int_q^{\frac{\theta - b}{\rho}} \theta - \rho x - b dx.$$

In the second case, $v^{-1}(b(q))$ equals Q, which holds for q close to Q (since these bids are close to 0). The conditional regret minimizing bid then solves

$$b \cdot q = \int_q^{Q} \theta - \rho x - b dx.$$

The cutoff $\bar{q}$ between the two cases is such that $v(Q) = b^{\text{UPA}}(\bar{q})$. Taken together, the conditional regret minimizing bid is

$$b^{\text{UPA}}(q) = \begin{cases} \theta - \sqrt{q\rho(2\theta - q\rho)} & \text{if } 0 \leq q \leq \frac{\theta - \sqrt{\theta^2 - \rho^2 Q^2}}{\rho} \\ \frac{(Q-q)(2\theta - \rho(q+Q))}{2Q} & \text{else.} \end{cases}$$

The bid is linear in q only if $\rho = 0$. In this case, $b^{\text{UPA}}(q) = \frac{\theta(Q-q)}{Q}$. ◀

5.3 Comparison of auction formats

We now compare the minimax-loss bids across auction formats. Previous theoretical work has identified uniform-price bids as more elastic (i.e., steeper) than pay-as-bid bids [Malvey and Archibald, 1998; Ausubel et al., 2014; Pycia and Woodward, 2025] in the Bayesian paradigm. This results from the significant demand-shading incentives for small quantities in the pay-as-bid auction—where bids for small quantities are paid for all larger quantities—and the

[29] The existence of collusive-seeming Bayes-Nash equilibria is linked to the absence of supply uncertainty [Klemperer and Meyer, 1989, Burkett and Woodward, 2020b]. See footnote 17 for an interpretation of supply uncertainty in our model.

significant demand-shading incentives for large quantities in the uniform-price auction—where bids are paid times the quantity for which they are offered. This intuition extends to the loss-minimization context, provided restriction is made to conditional regret-minimizing bids in the uniform-price auction. Define the average slope of the bid b to be $\alpha = (b(0) - b(Q))/Q$.

Comparison 1 (Uniform-price bids above pay-as-bid bids)**.** *The unique conditional regret-minimizing bid in the uniform-price auction is higher and on average steeper than the unique minimax-loss bid in the pay-as-bid auction:* $b^{UPA} \geq b^{PAB}$ *and* $\alpha^{UPA} \geq \alpha^{PAB}$.

Although conditional regret minimizing bids in the uniform-price auction are above the unique minimax-loss bid in the pay-as-bid auction, this is not the case for all selections of minimax-loss bids in the uniform-price auction. In the uniform-price auction, underbidding regret for large quantities is necessarily small: uniform pricing implies there is no wedge for overpayment (as there is in the pay-as-bid auction), and there is little utility foregone by not receiving a small number of units. Since conditional regret is the larger of overbidding and underbidding regret, and overbidding regret is increasing in bid, for large quantities there is a conditional regret-minimizing bid which is equal to zero; this zero bid is below the strictly positive minimax-loss bid in the pay-as-bid auction.

Comparison 2 (Semi-comparability of optimal bids)**.** *If* b^{UPA} *is a minimax-loss bid in the unconstrained uniform-price auction, then* $b^{UPA} \not\leq b^{PAB}$. *However, there is a minimax-loss bid* b^{UPA} *in the unconstrained uniform-price auction such that* $b^{UPA} \not\geq b^{PAB}$.

The comparisons of the bid functions imply that the auctioneer's revenues cannot be generally compared across the two auction formats. The uniform-price auction will lead to higher revenue when there are many similar bidders that all win a small quantity; this follows from Comparison 1. On the other hand, when the value distribution is such that a single bidder wins a large quantity, then price discrimination in the pay-as-bid auction will yield to higher revenue. These arguments establish the ambiguous revenue ranking.

Comparison 3 (Ambiguous revenue)**.** *Depending on the joint value distribution, both expected and ex post revenues can be higher in either auction format.*

While revenue cannot be ranked across auction formats, bidder loss is uniformly lower in the uniform-price auction than in the pay-as-bid auction. The existence of multiple minimax-loss bids in the uniform-price auction does not affect this comparison, because even when some bids are not uniquely defined, the level of minimax loss is.

Comparison 4 (Minimax loss)**.** *In the unconstrained case, minimax loss is lower in the uniform-price auction than in the pay-as-bid auction,*

$$\sup_{B^{-i}\in\mathcal{B}} L^{UPA}\left(b^{UPA};B^{-i},v^i\right) \leq \sup_{B^{-i}\in\mathcal{B}} L^{PAB}\left(b^{PAB};B^{-i},v^i\right).$$

What are the implications of one mechanism having lower minimax loss than another? Suppose a bidder can obtain costly information about the other bidders' behavior; this information will shrink the set of possible bid distributions $\mathcal{B}$. The bidder will tend to acquire more information when the subsequent auction mechanism yields higher minimax loss. Thus Comparison 4 implies that bidders in the pay-as-bid auction may obtain more costly information than bidders in the uniform-price auction.

5.4 Comparison to Bayes-Nash equilibrium

To situate our results in the literature, we compare bids under maximal uncertainty to those in Bayes-Nash equilibrium in the divisible-good framework of Ausubel et al. [2014]. In this model, bidders have symmetric linear marginal values $v(q;\theta) = (\theta - \rho q)_+$, $\rho > 0$, and aggregate supply is distributed according to a Pareto distribution with cumulative distribution function F given by $F(Q) = 1-(1+\xi Q/(\sigma n))^{-\frac{1}{\xi}}$, where ξ is the shape parameter and σn is the scale parameter. Equilibrium bids in the uniform-price and pay-as-bid auctions are $b^{\text{UPA}}_{\text{BNE}}$ and $b^{\text{PAB}}_{\text{BNE}}$, respectively, where

$$b^{\text{UPA}}_{\text{BNE}}(q) = \theta - \frac{n-1}{n-2}\rho q \quad \text{and} \quad b^{\text{PAB}}_{\text{BNE}}(q) = \theta - \frac{(n-1)\rho}{n(1-\xi)-1}\left(q + \frac{n}{n-1}\sigma\right).$$

The equilibrium in the uniform-price auction exists only if $n > 2$ and in the pay-as-bid auction only if $\xi < (n-1)/n$. For the pay-as-bid auction we assume that Q has bounded support, which is the case if $\xi < 0$. Let $\overline{Q}$ denote the upper bound of support. Equilibrium bids in the uniform-price auction constitute an ex-post equilibrium (cf. Klemperer and Meyer [1989]), while equilibrium bids in the pay-as-bid auction vary with the distribution of random supply. The maximum positive-value quantity $q^* = \theta/\rho$ will feature prominently in our analysis. Let $b^{\text{PAB}}_{\text{MML}}$ denote the unique minimax-loss bid in the pay-as-bid auction and $b^{\text{UPA}}_{\text{MML}}$ the unique conditional regret minimizing bid in the uniform-price auction.

5.4.1 Comparison of bids

In the uniform-price auction, bids under maximal uncertainty are below bids in BNE for small quantities, and may be above bids in BNE for large quantities. Bayesian bidders know that competitive pressure is high when there are many bidders or when marginal values are flat. This information leads bids to be close to value in the BNE while such information is unknown under maximal uncertainty.

Proposition 3 (Bids in UPA are often more aggressive in Bayes-Nash equilibrium)**.** *Either* $b^{UPA}_{BNE}(q) \geq b^{UPA}_{MML}(q)$ *for all* $q \in [0, \overline{Q}]$*, or there is a unique* $\hat{q}$ *such that* $b^{UPA}_{BNE}(q) > b^{UPA}_{MML}(q)$ *for all* $q \in (0, \hat{q})$ *and* $b^{UPA}_{BNE}(q) < b^{UPA}_{MML}(q)$ *for all* $q \in (\hat{q}, \overline{Q})$*. In either case, when* ρ *is sufficiently low, then* $b^{UPA}_{BNE}(q) > b^{UPA}_{MML}$ *for all* $q \in (0, \overline{Q})$*. As* n *becomes large,* $\hat{q} \to \overline{Q}$*.*

Importantly, equilibrium bids in the uniform-price auction under maximal uncertainty are unaffected by the number of bidders n. Since equilibrium bids become truthful in the Bayes-Nash equilibrium of the uniform-price auction as the number of bidders approaches infinity, in large markets the auctioneer may benefit from inducing BNE behavior where feasible;[30] of course, this will not be true if bidders play a tacitly collusive equilibrium as in the two-unit case.

With regard to the pay-as-bid auction, a general comparison is hindered by the complex dependency of BNE bids on the full distribution of aggregate quantity, while bids under maximal uncertainty depend only on the maximum quantity Q. Nonetheless, we observe the following.

Proposition 4 (Bids in PAB are often more aggressive in Bayes-Nash equilibrium)**.** *When* ρ *is sufficiently small, then* $b^{PAB}_{BNE}(q) > b^{PAB}_{MML}(q)$ *for all* $q < q^*$*. If* $\overline{Q}/n < q^*$*, then* $b^{PAB}_{BNE}(q) > b^{PAB}_{MML}(q)$ *for all* $q > \overline{Q}/n$*.*

As observed in Pycia and Woodward [2025], BNE bids will exceed values for unobtainable quantities (those above $\overline{Q}/n$). These bids are dominated by bidding value, implying that the BNE of the pay-as-bid auction is in dominated strategies. The minimax-loss bid is undominated.

5.4.2 Comparison of equilibrium outcomes

We now compare outcomes under maximal uncertainty to those in Bayes-Nash equilibrium. In a fixed context it is reasonable to expect strategies from only one of maximal uncertainty or BNE; comparison of equilibrium outcomes addresses the question of whether the auctioneer may benefit from driving behavior toward one solution concept or another. For example, if the auctioneer expects bidders to behave as if under maximal uncertainty, the auctioneer may be able to release credible information about the auction environment to steer bidders

[30] Although Proposition 3 is stated and proved in the model of Ausubel et al. [2014], the results of Swinkels [2001] imply that the Proposition generalizes to uniform-price auctions in general.

toward BNE strategies. We show below that this will generally be optimal in homogeneous environments.

We handle each auction format in turn. For pay-as-bid, recall that the supply-optimization results of Pycia and Woodward [2025] imply that when marginal values are linear, optimal revenue in the pay-as-bid auction is equal to half the area under the bidders' true demand curves.

Corollary 1 (Revenue in optimized pay-as-bid, Pycia and Woodward [2025])**.** *When bidders have symmetric linear marginal values, $v^i(q) = (\theta - \rho q)_+$, then the revenue generated in the Bayes-Nash equilibrium of an optimized pay-as-bid auction is $\pi^{PAB} = \theta/4\rho$.*

We now show that under maximal uncertainty even an optimized pay-as-bid auction generates less revenue than in the optimal Bayes-Nash auction. First, note that loss is increasing in aggregate quantity: for aggregate quantity Q either $v(Q) = 0$ and increasing supply has no impact on loss, or $v(Q) > 0$ and increasing supply, holding bids fixed, will increase loss for large quantities. Since Theorem 1 implies that, under maximum uncertainty, the total area under the bid curve is equal to loss, it follows that revenue is maximized when aggregate supply is as large as possible; henceforth we will take supply to be infinite. This implies the following revenue comparison.

Proposition 5 (Revenue comparison of optimized pay-as-bid)**.** *When aggregate supply is set to maximize revenue, Bayes-Nash equilibrium raises strictly greater expected revenue than maximal uncertainty.*

It is straightforward to establish Proposition 5 by graphical contradiction. Because marginal values are linear, optimal BNE per-capita revenue in the pay-as-bid auction is half the area under the bidder's true demand curve. In the pay-as-bid auction under maximal uncertainty, loss is constant across all possible allocations (Theorem 1). Loss at the maximum allocation is the area under the bid curve, which is the maximum obtainable revenue. Then if revenue is to be higher under maximal uncertainty than in BNE, loss must be at least half the area under the bidder's marginal value curve. Since loss is constant for all units, loss is equal to $\int_0^Q (v(q) - b(0))_+ dq$; that is, the initial bid $b(0)$ is set so that half of the area under the bidder's marginal value curve is above $b(0)$, and by implication half is below $b(0)$. But since bids are strictly decreasing where marginal values are positive and bounded above by $v(\cdot)$, it follows that the area under the bid curve is $\int_0^Q b(q)dq < \int_0^Q \min\{b(0), v(q)\}dq$. Then revenue under maximal uncertainty is less than half the area under the bidder's marginal value curve, and hence less than revenue in BNE.

Proposition 5 implies that if the auctioneer has sufficient information to optimize supply in a pay-as-bid auction, they should release this information if they may do so credibly. The following proposition implies that the same conclusion holds for the uniform-price auction.

Proposition 6 (Revenue comparison of optimized uniform-price)**.** *When aggregate supply is set to maximize revenue, Bayes-Nash equilibrium raises strictly greater expected revenue than maximum uncertainty.*

Maximal uncertainty is the revenue-optimal informational extreme when bidders play a collusive BNE in the uniform-price auction. In this case, the auctioneer's optimal information policy is not to reveal any information about past bidding.

6 Extension: Bidpoint-constrained minimax-loss bids

In practice bidders are frequently constrained from submitting a distinct bid for each quantity. For example, bidders can submit up to 10 bidpoints in Czech treasury auctions [Kastl, 2011] or 40 steps in the Texas electricity market [Hortaçsu et al., 2019]. We now consider the case in which bidder i can submit up to M bid points, $\{(q_{ik}, b_{ik})\}_{k=1}^{M}$, where $q_{ik} \leq q_{ik+1}$ and $b_{ik} \geq b_{ik+1}$ for all k. The implied bid function is a step function

$$\hat{b}^i(q) = \begin{cases} b_{ik} & \text{if } q_{k-1} \leq q < q_k, \\ 0 & \text{if } q = Q, \end{cases}$$

where $q_0 = 0$. Importantly, the quantities at which bids are submitted are a choice variable for the bidder. In this section we summarize the findings; detailed arguments are found in Appendix C.

6.1 Pay-as-bid auctions

As in the unconstrained case, the minimax-loss bid in the bidpoint-constrained pay-as-bid auction equates underbidding regret across all units. The minimax-loss bid is then found by solving a constrained optimization problem. Intuitively, the bidder minimizes their maximum payment subject to equal conditional regret across all outcomes. We illustrate the bidding function for the case in which the bidder has flat marginal values.

Example 4 (Pay-as-bid with flat marginal values)*.* Suppose bidder i's marginal value is flat,

$v^i(q) = \theta$ for all q. The constrained loss optimization problem is

$$\min_{q',b'} (\theta - b_1') Q, \text{ s.t. } \left(Q - q_{k-1}'\right)(\theta - b_k') + \sum_{k'=1}^{k} \left(q_{k'}' - q_{k'-1}'\right)(b_{k'}' - b_k') = (\theta - b_1') Q.$$

Equating conditional loss across units requires $R_{q_{k+1}}^{\text{PAB}} - R_{q_k}^{\text{PAB}} = 0$, or

$$0 = -Qb_{k+2} - (q_{k+1} - q_k)\theta + (Q + (q_{k+1} - q_k))\, b_{k+1}.$$

Solving this equation recursively, backwards from $b_{M+1} = 0$, gives a closed-form expression for optimal bids conditional on quantities,

$$b_k = \sum_{k'=k}^{M} \frac{Q^{k'-k}(q_{k'} - q_{k'-1})}{\prod_{j=k}^{k'} [Q + (q_j - q_{j-1})]} \theta.$$

Minimizing loss then implies

$$q_k = \frac{k}{M} Q, \text{ and } b_k = \frac{\theta}{M+1} \sum_{k'=k}^{M} \left[\frac{M}{M+1}\right]^{k'-k}.$$

Notably, minimax bidpoints are evenly spaced in the quantity space. Figure 5 plots these bids and compares them to minimax-loss bids of the corresponding uniform-price auction.◀

6.2 Uniform-price auctions

In Section 5.2 we showed that there are typically many minimax-loss bids in the unconstrained uniform-price auction. We show in Appendix C that this is in stark contrast to the bidpoint-constrained uniform-price auction, where there is a unique minimax-loss bid.

We provide some intuition for the uniqueness in the constrained case and contrast it with the multiplicity of the unconstrained case. Intuitively, when the bidder receives a small quantity, they do not leave a lot of money on the table due to overbidding, because they received a small number of units and their total payment is low; they also do not miss out on significant utility from underbidding, because the market price will tend to be high and they will not desire many units at this price. Thus the main source of loss is bids on intermediate quantities, leaving bids on small (and very large) quantities only partially specified. This stands in contrast to the bidpoint-constrained case where the locations of the bid steps are choice variables. Given the choice, the bidder will submit relatively dense bids for intermediate quantities and relatively sparse bids for extreme quantities; the large gaps

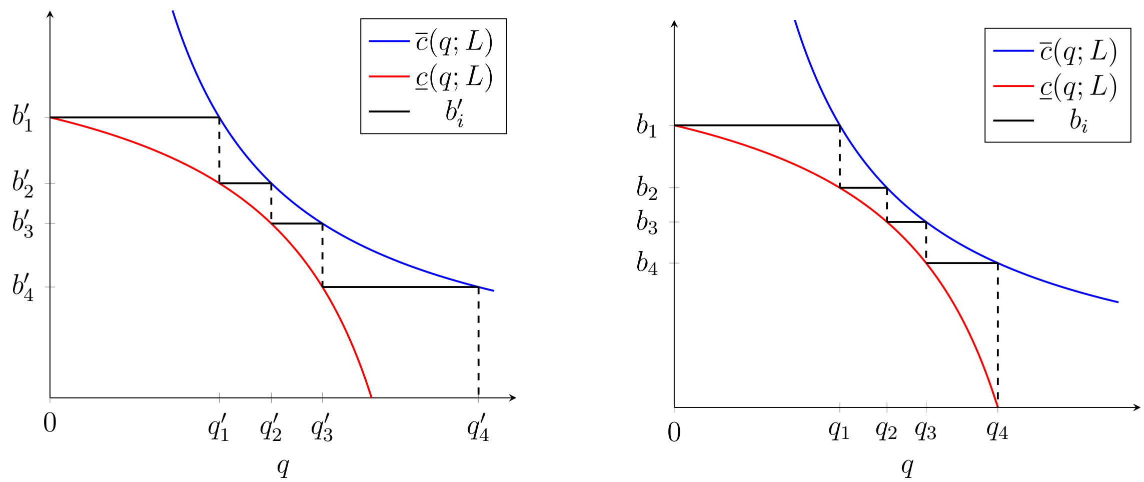

(a) Loss L above constrained minimax loss (b) Loss L equal to constrained minimax loss

Figure 4: Iso-loss curves of conditional underbidding and overbidding regret in the uniform-price auction.

between bid points for small units work against the intuition arising from the multi-unit case, where bidpoint gaps are uniform, that bids for small quantities are not uniquely determined.

The construction of the minimax-loss bid in the constrained uniform-price auction follows from observing that steps in the implied bid function extend between the two iso-loss curves. In particular, the minimax-loss bid in the constrained uniform-price auction extends from the lower iso-loss curve to the upper iso-loss curve, then jumps down to the lower iso-loss curve, and extends again to the upper iso-loss curve; this continues until a bid of zero is reached. Figure 4b illustrates this construction for $M = 4$. If the bid did not extend fully between the two iso-loss curves, with a slight perturbation the bid could be made to lie strictly between the two iso-loss curves, which would entail strictly lower loss. Constructing bidpoint-constrained minimax-loss bids is then straightforward. For loss L such that $\bar{c}(\cdot; L) \geq \underline{c}(\cdot; L)$, let $q_0 = 0$ and for all $k \in \{1, \ldots, M\}$ let $b_k = \underline{c}(q_{k-1}; L)$ and let q_k be such that $\bar{c}(q_k; L) = b_k$.[31] If $\underline{c}(q_M; L) > 0$ constrained minimax loss is above L, and if $\underline{c}(q_M; L) < 0$ constrained minimax loss is below L. In either case, a new level of loss L' may be proposed, and the procedure continues until $\underline{c}(q_M; L) = 0$. Figure 4a illustrates the case when the level of loss is above the minimax loss. In the figure, the final step q_4' is too high, and loss can be decreased.

The construction of minimax-loss bids between the upper and lower iso-loss curves provides an intuitive argument for the uniqueness of minimax-loss bids in the uniform-price auction. Given a level of loss and associated iso-loss curves, either there is no M-step step function between them, or there is a single M-step step function between them, or there are

[31] In the event that $\bar{c}(Q; L) > b_k$, we define $q_k = Q$.

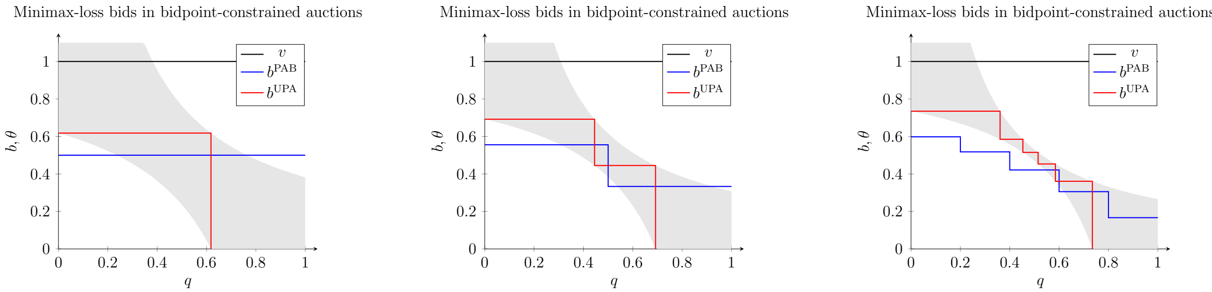

Figure 5: Minimax-loss bids under flat marginal values, when bidders are constrained to $M \in \{1, 2, 5\}$ bidpoints. In bidpoint-constrained auctions, minimax-loss bids are unique in both the pay-as-bid and uniform-price auctions. As the number of bidpoints increases, the upper and lower iso-loss curves in the uniform-price auction approach tagency.

multiple such step functions between them. If there is no feasible step function between the iso-loss curves, this level of loss is not feasible and minimax loss is above the assumed loss. On the other hand, if there are multiple feasible step functions between the iso-loss curves the iso-loss curves can be brought closer together (by reducing assumed loss) while still allowing for a feasible step function between them. This improvement in loss is infeasible only when there is a unique step function between the iso-loss curves, and at that point maximum loss is minimized.

The following example illustrates the constrained minimax-loss bid function when the bidder has flat marginal values.

Example 5 (Uniform-price with flat marginal values). Suppose that bidder i's marginal value v^i is flat, $v^i(q) = \theta$ for all q. The constrained loss optimization problem is

$$\min_{q',b'} b_1' q_1', \text{ s.t. } b_k' q_k' = (\theta - b_k') (Q - q_k') \quad \forall k.$$

The minimax-loss bid induces loss $C_M Q\theta$, and solves

$$q_0 = 0, \quad q_k = \left(C_M - \frac{C_M^2}{q_{k-1} - (1 - C_M)} \right) Q, \quad q_M = (1 - C_M) Q, \quad \text{and } b_k = \frac{C_M \theta}{q_k}.$$

The solution to this expression is unique: the recursive equation for q_k increases in C_M, while the endpoint condition for q_M decreases in C_M.[32] Figure 5 illustrates these bids and compares them to the unique minimax-loss bids in the pay-as-bid auction. ◀

Examples 4 and 5 suggest a new testable prediction. With flat marginal values, the bids in the bidpoint-constrained pay-as-bid auction are evenly spaced, while they are more

[32]In the unconstrained model (Section 5.2) minimax loss is $\theta Q/4$. Since loss is higher when bids are constrained than when they are unconstrained, it follows that $q_1 \geq Q/3$ and $q_M \leq 3Q/4$. That is, minimax-loss bid points are all for interior quantities.

clustered around intermediate quantities in the bidpoint-constrained uniform-price auction. More generally, the location of the bids in the pay-as-bid auction is more dispersed than in the uniform-price auction.

The examples also show that Comparison 1, which shows that the conditional regret minimizing bid of the uniform-price auction is higher and steeper than the minimax-loss bid of the pay-as-bid auction, does not fully extend to the bidpoint-constrained case. In the constrained case, bids are on average steeper in the uniform-price auction than in the pay-as-bid auction, $\alpha^{\text{UPA}} \geq \alpha^{\text{PAB}}$, but neither auction's bids are higher: $b_1^{\text{UPA}} > b_1^{\text{PAB}}$ and $q_M^{\text{UPA}} < q_M^{\text{PAB}}$. By continuity, even if the marginal values are not perfectly flat, the two bid functions cannot be ranked uniformly in the constrained case.

7 Conclusion

In this paper we have characterized optimal prior-free bids in the pay-as-bid and uniform-price auctions, the two leading auction formats for allocating homogeneous goods such as electricity and government debt. The two pricing rules create different incentives for the bidders; our analysis shows that taking a worst-case loss approach to bid optimization enables a tractable analysis of the two formats and leads to new testable predictions. Remarkably, our analysis remains tractable even with multi-dimensional private information because we do not require the inversion of strategies as in the canonical Bayes-Nash equilibrium approach. Hence, we believe the worst-case loss approach may also be fruitfully applied to other complex strategic interactions.

References

Robert Aumann and Adam Brandenburger. Epistemic conditions for nash equilibrium. *Econometrica*, 63(5):1161–1180, 1995.

Lawrence M Ausubel, Peter Cramton, Marek Pycia, Marzena Rostek, and Marek Weretka. Demand reduction and inefficiency in multi-unit auctions. *The Review of Economic Studies*, 81(4):1366–1400, 2014.

Kerry Back and Jaime F Zender. Auctions of divisible goods: on the rationale for the treasury experiment. *Review of Financial Studies*, 6(4):733–764, 1993.

Klenio Barbosa, Dakshina G. De Silva, Liyu Yang, and Hisayuki Yoshimoto. Auction mechanisms and treasury revenue: Evidence from the chinese experiment. *American Economic Journal: Microeconomics*, 14(4):394–419, November 2022.

Dirk Bergemann and Karl H. Schlag. Pricing without priors. *Journal of the European Economic Association*, 6(2-3):560–569, 2008.

Dirk Bergemann and Karl H. Schlag. Robust monopoly pricing. *Journal of Economic Theory*, 146(6):2527–2543, 2011.

Omar Besbes and Assaf Zeevi. On the minimax complexity of pricing in a changing environment. *Operations Research*, 59(1):66–79, 2011.

Daniel Bougt, Gagan Ghosh, and Heng Liu. Revenue effects of ambiguity in multi-unit auctions. 2024.

Menachem Brenner, Dan Galai, and Orly Sade. Sovereign debt auctions: Uniform or discriminatory? *Journal of Monetary Economics*, 56(2):267–274, 2009.

Justin Burkett and Kyle Woodward. Uniform price auctions with a last accepted bid pricing rule. *Journal of Economic Theory*, 2020a.

Justin Burkett and Kyle Woodward. Reserve prices eliminate low revenue equilibria in uniform price auctions. *Games and Economic Behavior*, 2020b.

Vincent P. Crawford. New directions for modelling strategic behavior: Game-theoretic models of communication, coordination, and cooperation in economic relationships. *Journal of Economic Perspectives*, 30(4):131–50, November 2016.

Pablo Del Río. Designing auctions for renewable electricity support. best practices from around the world. *Energy for Sustainable Development*, 41:1–13, 2017.

Ulrich Doraszelski, Gregory Lewis, and Ariel Pakes. Just starting out: Learning and equilibrium in a new market. *American Economic Review*, 108(3):565–615, 2018.

Richard Engelbrecht-Wiggans and Charles M. Kahn. Multi-unit auctions with uniform prices. *Economic Theory*, 12(2):227–258, 1998.

Richard Engelbrecht-Wiggans and Charles M Kahn. Multiunit auctions in which almost every bid wins. *Southern Economic Journal*, pages 617–631, 2002.

Dirk Engelmann and Veronika Grimm. Bidding Behaviour in Multi-Unit Auctions – an Experimental Investigation. *The Economic Journal*, 119(537):855–882, 03 2009.

Yingni Guo and Eran Shmaya. Robust monopoly regulation. *American Economic Review*, 115(2):599–634, 2025.

Yingni Guo and Eran Shmaya. Regret-minimizing project choice. *Econometrica*, 91:1567–1593, 2023.

Joseph Y. Halpern and Rafael Pass. Iterated regret minimization: A new solution concept. *Games and Economic Behavior*, 74(1):184 – 207, 2012.

C.-Philipp Heller and Lorenz Wieshammer. Die *Merit Order* und Auktionen auf den europäischen Strommärkten: Marktdesign als Preiskartell? *Neue Zeitschrift für Kartellrecht*, 2023.

Ali Hortaçsu and Jakub Kastl. Valuing dealers' informational advantage: A study of canadian treasury auctions. *Econometrica*, 80(6):2511–2542, November 2012.

Ali Hortacsu and David McAdams. Mechanism choice and strategic bidding in divisible good auctions: an empirical analysis of the turkish treasury auction market. *Journal of Political Economy*, 118(5):833–865, 2010.

Ali Hortaçsu, Fernando Luco, Steven L. Puller, and Dongni Zhu. Does strategic ability affect efficiency? evidence from electricity markets. *American Economic Review*, 109(12): 4302–42, December 2019.

Bernhard Kasberger. An equilibrium model of the first-price auction with strategic uncertainty: Theory and empirics. *Working paper*, 2022.

Bernhard Kasberger and Karl H. Schlag. Robust bidding in first-price auctions: How to bid without knowing what others are doing. *Management Science*, 70(7):4219–4235, 2024.

Bernhard Kasberger and Kyle Woodward. Bidding in multi-unit auctions under limited information. *arXiv preprint: 2112.11320*, 2023.

Jakub Kastl. Discrete Bids and Empirical Inference in Divisible Good Auctions. *The Review of Economic Studies*, 78(3):974–1014, 02 2011.

Jakub Kastl. On the properties of equilibria in private value divisible good auctions with constrained bidding. *Journal of Mathematical Economics*, 48(6):339–352, 2012.

Paul D Klemperer and Margaret A Meyer. Supply function equilibria in oligopoly under uncertainty. *Econometrica: Journal of the Econometric Society*, pages 1243–1277, 1989.

P.B. Linhart and R. Radner. Minimax-regret strategies for bargaining over several variables. *Journal of Economic Theory*, 48(1):152–178, 1989.

Paul F Malvey and Christine M Archibald. Uniform-price auctions: Update of the treasury experience. *US Treasury*, 1998.

Charles F. Manski. Econometrics for decision making: Building foundations sketched by haavelmo and wald. *Econometrica*, 89(6):2827–2853, 2021.

Daniel Marszalec, Alexander Teytelboym, and Sanna Laksá. Epic fail: How below-bid pricing backfires in multiunit auctions. *CIRJE Working Paper Series,, CIRJE-F-1096*, 2020.

Luiz Maurer and Luiz A Barroso. *Electricity auctions: an overview of efficient practices.* World Bank Publications, 2011.

David McAdams. On the failure of monotonicity in uniform-price auctions. *Journal of Economic Theory*, 137(1):729–732, 2007.

Charles Noussair. Equilibria in a multi-object uniform price sealed bid auction with multi-unit demands. *Economic Theory*, 5(2):337–351, 1995.

OECD. *OECD Sovereign Borrowing Outlook 2021.* 2021. https://doi.org/10.1787/48828791-en.

Alexei Parakhonyak and Anton Sobolev. Non-reservation price equilibrium and search without priors. *Economic Journal*, 125:887–909, 2015.

Georgia Perakis and Guillaume Roels. Regret in the newsvendor model with partial information. *Operations Research*, 56(1):188–203, 2008.

Marek Pycia and Kyle Woodward. A case for pay-as-bid auctions. 2025.

Itzhak Rasooly. Going. . . going. . . wrong: A test of the level-k (and cognitive-hierarchy) models of bidding behavior. *Journal of Political Economy Microeconomics*, 1(2):400–445, 2023.

Ludovic Renou and Karl H. Schlag. Minimax regret and strategic uncertainty. *Journal of Economic Theory*, 145(1):264 – 286, 2010.

L. J. Savage. The theory of statistical decision. *Journal of the American Statistical Association*, 46(253):55–67, 1951.

Karl H. Schlag and Andriy Zapechelnyuk. Robust sequential search. *Theoretical Economics*, 16(4):1431 – 1470, 2021.

Karl H. Schlag and Andriy Zapechelnyuk. Compromise, don't optimize: Generalizing perfect bayesian equilibrium to allow for ambiguity. *Journal of Political Economy Microeconomics*, 2(1):77–128, 2023.

Jörg Stoye. Axioms for minimax regret choice correspondences. *Journal of Economic Theory*, 146(6):2226 – 2251, 2011.

Jeroen M Swinkels. Efficiency of large private value auctions. *Econometrica*, 69(1):37–68, 2001.

James JD Wang and Jaime F Zender. Auctioning divisible goods. *Economic Theory*, 19(4): 673–705, 2002.

Gerhard Winkler. Extreme points of moment sets. *Mathematics of Operations Research*, 13 (4):581–587, 1988.

Kyle Woodward. Strategic ironing in pay-as-bid auctions: Equilibrium existence with private information. Working paper, 2016.

Kyle Woodward. Mixed-price auctions for divisible goods. 2021.

A Omitted Proofs

A.1 Proofs for Section 3

Proof of Proposition 1. Because there are only two units for sale, it is sufficient to analyze the two bidders with the highest marginal values for their first units. Denote these bidders 1 and 2, and let bidder i's marginal values be $(v_i, \rho v_i)$. Without loss of generality assume that bidder 1's marginal value for their first unit is above bidder 2's, $v_1 > v_2$. Let $\tau \in [0,1]$ be such that $v_2 = \tau v_1$. Suppose it is efficient that the two bidders with the highest values win one unit each. This can only be the case if $v_1 + \tau v_1 \geq v_1 + \rho v_1$, i.e., if and only if $\tau \geq \rho$. Note that it cannot be efficient that the bidder with the third-highest value wins anything.

We first show the result when having two winners is efficient. Let $\tau \geq \rho$. Bidder 1 wins one unit in the PAB auction since $b_{21}^{\text{PAB}} \geq b_{12}^{\text{PAB}}$. If $\rho \leq \frac{3}{7}$, then this inequality equals $\frac{\tau v_1(3-\rho)}{6} \geq \frac{\rho v_1}{3}$, which is true by assumption. If $\rho > \frac{3}{7}$, then $\frac{\tau v_1(3+2\rho)}{9} \geq \frac{\rho v_1}{3}$ is also true because $\tau \geq \rho \geq \frac{3\rho}{3+2\rho}$ holds.

Bidder 1 wins one unit in the FRB uniform-price auction since $b_{21}^{\text{FRB}} \geq b_{12}^{\text{FRB}} \Leftrightarrow \tau v_1 \geq \rho v_1/2 \Leftrightarrow \tau > \rho/2$, which is true by assumption.

Bidder 1 wins one unit in the LAB uniform-price auction since $b_{21}^{\text{LAB}} \geq b_{12}^{\text{LAB}}$ holds. To see this, note that the inequality is equivalent to $\frac{\tau v_1}{2} \geq \frac{\rho v_1}{3}$ if $\rho < \frac{1}{2}$. This inequality holds by assumption. If $\rho \geq \frac{1}{2}$, then bidder 1 wins one unit if $\frac{\tau v_1(1+\rho)}{3} \geq \frac{\rho v_1}{3}$, which also holds since $\tau \geq \rho \geq \frac{\rho}{1+\rho}$.

We now prove the welfare ranking when it is *ex post* efficient for a single bidder to receive both units. Let $\rho > \tau$. Let $\rho \leq \frac{3}{7}$. From the previous paragraphs, bidder 1 wins two units in the PAB auction if $\tau \leq \frac{2\rho}{3-\rho}$. Bidder 1 wins two units in the FRB auction if $\tau \leq \frac{\rho}{2}$. Bidder 1 wins two units in the LAB auction if $\tau \leq \frac{2\rho}{3}$. Since the PAB cutoff for τ is higher than the LAB cutoff and the LAB cutoff is higher than the FRB cutoff ($\frac{2\rho}{3-\rho} \geq \frac{2\rho}{3} \geq \frac{\rho}{2}$), the PAB is efficient whenever the LAB is efficient and the LAB is efficient whenever the FRB is efficient. The same ranking of cutoffs applies when $\frac{3}{7} \leq \rho \leq \frac{1}{2}$ (in which case it is $\frac{3\rho}{3+2\rho} \geq \frac{2\rho}{3} \geq \frac{\rho}{2}$) and when $\rho > \frac{1}{2}$ ($\frac{3\rho}{3+2\rho} \geq \frac{\rho}{1+\rho} \geq \frac{\rho}{2}$). □

A.2 Proofs for Section 4

Proof of Lemma 1. Consider the maximization of loss

$$\sup_{\tilde{b}} \sup_{B^{-i} \in \mathcal{B}} \mathbb{E}_{B^{-i}} \left[\hat{u}\left(q^i(\tilde{b}, b^{-i}), t^i\left(\tilde{b}, b^{-i}\right); v^i\right) - \hat{u}\left(q^i\left(b^i, b^{-i}\right), t^i\left(b^i, b^{-i}\right); v^i\right)\right],$$

where we have swapped the order of the suprema. Observe that the inner maximization problem is linear in the choice variable B^{-i}. Winkler [1988] proves that the extreme points of $\mathcal{B}$ are distributions with a single point in the support. Since loss is linear in B^{-i}, maximum loss is attained at an extreme point. □

A.2.1 Analysis of pay-as-bid auctions

Proof of Lemma 2. Lemma 1 proves that loss is maximized by bid distributions with one opponent bid profile in the support, leading to the equivalence of loss and regret. Consider regret $R\left(b^i; b^{-i}, v^i\right)$ and suppose b^{-i} is such that bidder i wins q units: $q = q^i(b^i, b^{-i})$. Regret depends on bidder i's best reply to b^{-i}, which is given by $\tilde{b}^i \in \arg\sup_{\tilde{b}} \hat{u}\left(q^i\left(\tilde{b}, b^{-i}\right), t^i\left(\tilde{b}, b^{-i}\right); v^i\right)$. Let $q' = q^i(\tilde{b}^i, b^{-i})$. Regret is then

$$R\left(b^i; b^{-i}, v^i\right) = \int_0^{q'} v^i(x) - \tilde{b}^i(x)dx - \int_0^q v^i(x) - b^i(x)dx.$$

Regret is decreasing pointwise in $\tilde{b}^i$. Instead of maximizing regret with respect to b^{-i}, we maximize it with respect to $\tilde{b}^i$ and q'. Note that q' cannot be strictly lower than q since $b^i(x) \leq v^i(x)$ for all x.

If $q' > q$, then regret can be written as

$$\int_0^q b^i(x) - \tilde{b}^i(x)dx + \int_q^{q'} v^i(x) - \tilde{b}^i(x)dx.$$

To win q' units, $\tilde{b}^i(x) \geq b^i(q)$ must be true for all $x \in [0, q']$. To see this, note that bidder i does not win more than q units with bid $b^i(q)$. Hence, to win $q' > q$ units, bidder i needs to bid at least this much. The quantity q' that then maximizes regret is either Q or such that $v^i(q') = b^i(q)$. This leads to the expression for underbidding regret in Equation (6). A worst-case opponent bid profile is such that they all submit flat bids at $b^i(q)$.

If $q' = q$, then regret equals

$$\int_0^q b^i(x) - \tilde{b}^i(x)dx.$$

Regret is clearly maximized if $\tilde{b}^i(x) = 0$, leading to overbidding regret $\overline{R}_q^{\mathrm{PAB}}(b^i; v^i)$. A worst-case bid profile is such that there is one other bidder who bids $v^i(0)$ for quantities below $Q - q$ and nothing else. □

A.2.2 Analysis of uniform-price auctions

Proof of Lemma 3. The proof of this claim is substantially similar to the proof of the equivalent result for the pay-as-bid auction (Lemma 2) and is omitted. □

A.3 Proofs for Section 5

A.3.1 Pay-as-bid auctions

Proof of Lemma 4. We show that $\underline{R}_q^{\mathrm{PAB}}(b^i; v^i) = \int_0^Q b^i(x)dx$ for all q. First, since $\int_0^q b^i(x)dx$ is weakly increasing in q, Lemma 2 implies that maximum loss is

$$\max\left\{\sup_{q\in[0,Q)} \underline{R}_q^{\mathrm{PAB}}\left(b^i; v^i\right), \int_0^Q b^i(x)dx\right\}.$$

Note that increasing all bids by $\varepsilon > 0$ will weakly decrease $\underline{R}_q(b^i; v^i)$ for all q and strictly increase $\int_0^Q b^i(x)dx$. Then, if b^i is loss-minimizing, it must be that $\int_0^Q b^i(x)dx \geq \sup_q \underline{R}_q^{\mathrm{PAB}}(b^i; v^i)$. Similarly, decreasing all bids by $\varepsilon > 0$ strictly decreases $\int_0^Q b^i(x)dx$ and continuously affects

$\underline{R}_q^{\text{PAB}}(b^i; v^i)$, thus $\int_0^Q b^i(x)dx = \sup_q \underline{R}_q^{\text{PAB}}(b^i; v^i)$.[33]

Now, suppose that there is $q \in [0, Q)$ with $\underline{R}_q^{\text{PAB}}(b^i; v^i) < \int_0^Q b^i(x)dx$. If $b^i(q) = 0$, then

$$\underline{R}_q^{\text{PAB}}\left(b^i; v^i\right) = \int_0^q b^i(x)dx + \int_q^Q v^i(x)dx \geq \int_0^Q b^i(x)dx = \int_0^q b^i(x)dx.$$

This is a contradiction, and it must be that $b^i(q) > 0$. In this case, reducing $b^i(q)$ will weakly increase $\underline{R}_q^{\text{PAB}}(b^i; v^i)$, strictly decrease $\underline{R}_{q'}^{\text{PAB}}(b^i; v^i)$ for all $q' > q$, and will not affect $\underline{R}_{q'}^{\text{PAB}}(b^i; v^i)$ for $q' < q$; reducing $b^i(q)$ also reduces $\int_0^Q b^i(x)dx$, and the arguments above show that increasing all bids by some small amount will strictly reduce loss. It follows that $\underline{R}_q^{\text{PAB}}(b^i; v^i) = \int_0^Q b^i(x)dx$ for all q. □

Proof of Theorem 1. Lemma 4 establishes that the derivative of underbidding regret must equal zero for all $q \in [0, Q]$. Recall that underbidding regret can be written as

$$\int_0^q b(x) - b(q)dx + \int_q^{v^{-1}(b(q))} v(x) - b(q)dx.$$

The first derivative of underbidding regret with respect to q is

$$-\int_0^q b'(q)dx - \int_q^{v^{-1}(b(q))} b'(q)dx + (v(v^{-1}(b(q))) - b(q))\frac{1}{v'(v^{-1}(b(q)))} - (v(q) - b(q)).$$

This straightforwardly simplifies to the derivative set equal to zero in Equation (7).

It remains to establish the initial condition and uniqueness. Because $b^i(Q) \geq 0$ by constraint, it is sufficient to show that $b^i(Q)$ cannot be strictly positive. By the fundamental theorem of differential equations (the Picard–Lindelöf theorem), if there are solutions b^i and $\tilde{b}^i$ with $b^i(Q) = 0 < \tilde{b}^i(Q)$, then $b^i \leq \tilde{b}^i$. The differential form ensures equal conditional regret for all units, and conditional regret for unit $q = Q$ under bid $\tilde{b}^i$ is $\int_0^Q \tilde{b}^i(x)dx > \int_0^Q b^i(x)dx$. Then maximum loss is lower under bid b^i than under bid $\tilde{b}^i$, and $\tilde{b}^i$ is not a minimax-loss bid. Then $b^i(Q) = 0$ for any minimax-loss bid, and uniqueness follows from the fundamental theorem of differential equations.

We now show that the minimax-loss bid vector is strictly below marginal values wherever $v > 0$. To see this, recall that we assumed $v(Q) > 0$ (and the corresponding discussion in footnote 10). The above implies that $b^{\text{PAB}}(Q) = 0 < v(Q)$. Then if there is q with $b^{\text{PAB}}(q) = v(q) > 0$, there is a maximal such quantity (because db^{PAB}/dq is continuous),

[33] A bid function which is not strictly positive—i.e., for which there exists q with $b^i(q) = 0$—cannot be uniformly decreased by ε. Nonetheless, decreasing the bid by ε where possible will decrease $\int_0^Q b^i(x)dx$ and will continuously affect $\sup_q \underline{R}_q^{\text{PAB}}(b^i; v^i)$.

denoted by $\bar{q}$. Equation (7) implies that the derivative of b^{PAB} is zero at $\bar{q}$, which contradicts the assumption that $\bar{q}$ is the maximal quantity with $b^{\text{PAB}}(q) = v(q) > 0$. It follows that $b^{\text{PAB}}(q) < v(q)$ for all $q \in [0, Q]$.

The minimax-loss bid vector is strictly decreasing in quantity wherever $v > 0$. This follows from the left-hand side of Equation (7) being negative (due bids being below marginal values) and inverse marginal values (v^{-1}) being positive.

If $v(Q) \geq b(0)$, then $v^{-1}(b(q)) = Q$ for all $q \in [0, Q]$. It can be directly verified that the first derivative of the bidding function in Equation (8) is as in Equation (7). $\square$

A.3.2 Uniform-price auctions

Proof of Proposition 2. Note that maximal regret is strictly positive for any non-degenerate v^i. Consider the bid $b^i(0)$. Overbidding regret is $\overline{R}_0^{\text{UPA}}(b^i; v^i) = 0 \cdot b^i(0) = 0$ and underbidding regret equals $\underline{R}_0^{\text{UPA}}(b^i; v^i) = \int_0^Q \left(v^i(x) - b^i(0)\right)_+ dx$. Hence, since minimax regret is positive, any bid $b^i(0) = v^i(0) - \varepsilon$ with $\varepsilon > 0$ but small is optimal. $\square$

A.3.3 Comparison of auction formats

Proof of Comparison 1. We first prove that when b^{UPA} is conditionally regret minimizing and b^{PAB} minimizes loss in the pay-as-bid auction, then $b^{\text{UPA}}(q) \geq b^{\text{PAB}}(q)$ for all $q \in [0, Q]$. Note first that $b^{\text{PAB}}(Q) = b^{\text{UPA}}(Q) = 0$. Next, observe that if $b^{\text{PAB}}(q) = b^{\text{UPA}}(q) = b$, then

$$\frac{db^{\text{PAB}}(q)}{dq} = -\frac{v(q) - b}{v^{-1}(b)} > -\frac{v(q)}{v^{-1}(b)} = \frac{db^{\text{UPA}}(q)}{dq}. \tag{9}$$

Hence, if the bids for a quantity are the same, then the absolute value of the slope of the UPA bid function is higher than the absolute value of the slope of the PAB bid function. Consequently, for q' marginally below q, the UPA bids are strictly higher than the PAB bids. It follows that $b^{\text{PAB}}(q') < b^{\text{UPA}}(q')$ for $q' < Q$ as well as the overall comparison.

The comparison of the average slope follows immediately from $b^{\text{PAB}}(Q) = b^{\text{UPA}}(Q) = 0$ and $b^{\text{PAB}}(0) < b^{\text{UPA}}(0) = v(0)$. $\square$

Proof of Comparison 2. In light of Comparison 4, we first show that the initial bid in the uniform-price auction must lie above the initial bid in the pay-as-bid auction. Specifically, $R_0^{\text{PAB}}(b^{\text{PAB}}; v^i) = \int_0^Q (v^i(x) - b^{\text{PAB}}(0))_+ dx = L^{\text{PAB}}$ and $R_0^{\text{UPA}}(b^{\text{UPA}}; v^i) = \int_0^Q (v^i(x) - b^{\text{UPA}}(0))_+ dx \leq L^{\text{UPA}} \leq L^{\text{PAB}}$. It follows that $b^{\text{UPA}}(0) \geq b^{\text{PAB}}(0)$, and thus it cannot be that $b^{\text{UPA}} < b^{\text{PAB}}$.

Second, observe that for q close to Q underbidding loss becomes arbitrarily close to 0 in the uniform-price auction. Thus, the lower iso-loss curve must intersect the horizontal axis at some $q < Q$, implying the existence of a minimax-loss bid which is zero for quantities strictly below Q. Since the minimax-loss bids in the pay-as-bid auction are positive for all $q < Q$ (Theorem 1), there exists a minimax-loss bid in the uniform-price auction which is not everywhere above the unique minimax-loss bid in the unconstrained pay-as-bid auction. $\square$

Proof of Comparison 4. Let q be the quantity for which worst-case loss equals conditional regret in the uniform-price auction, and let b^{UPA} denote the conditional regret minimizing bids of the uniform-price auction. Then we have that

$$\begin{aligned}\sup_{B^{-i}\in\mathcal{B}} L^{\text{UPA}}\left(b^{\text{UPA}}; B^{-i}, v^i\right) &= \int_q^Q \left(v^i\left(x\right) - b^{\text{LAB}}\left(q\right)\right)_+ dx \\ &\leq \int_q^Q \left(v^i\left(x\right) - b^{\text{PAB}}\left(q\right)\right)_+ dx \\ &\leq \int_q^Q \left(v^i\left(x\right) - b^{\text{PAB}}\left(q\right)\right)_+ dx + \int_0^q b^{\text{PAB}}(x) - b^{\text{PAB}}(q)dx \\ &= \sup_{B^{-i}\in\mathcal{B}} L^{\text{PAB}}\left(b^{\text{PAB}}; B^{-i}, v^i\right),\end{aligned}$$

where we use that $b^{\text{PAB}} \leq b^{\text{UPA}}$ (Comparison 1) and the fact that underbidding regret involves lowering the bids on $[0, q]$. $\square$

A.3.4 Comparison of bids

From Example 3, define functions b_L and b_R by

$$b_L\left(q\right) = \theta - \sqrt{\left(2\theta - \rho q\right)\rho q}, \quad b_R\left(q; Q\right) = \frac{\left(Q - q\right)\left(2\theta - \left(q + Q\right)\rho\right)}{2Q}.$$

With these definitions, we can write the conditional regret minimizing bid in the uniform-price auction as

$$b^{\text{UPA}}\left(q\right) = \begin{cases} b_L\left(q\right) & \text{if } 0 \leq q \leq \frac{\theta - \sqrt{\theta^2 - \rho^2 Q^2}}{\rho}, \\ b_R\left(q\right) & \text{otherwise.}\end{cases}$$

Lemma 5 (Comparison of piecewise components of b^{UPA})**.** *For all* $q \in [0, \overline{Q}]$, $b_L(q) \geq b_R(q)$.

Proof. The bid function b_L equates overbidding and underbidding loss for quantity q when $b(q) > v(Q)$, and the bid function b_R equates overbidding and underbidding loss for quantity q when $v^{-1}(b(q)) = Q$. Note that Q does not directly affect b_L. Moreover, whenever $v(Q) > 0$, increasing Q increases the *underbidding* loss associated with any quantity q such that $v(Q) > b_R(q; Q)$ without affecting the overbidding loss. Then increasing Q must weakly increase bids for such quantities. For any such quantity bids increase in Q until $b_R(q;Q) = b_L(q)$, and beyond this point the bid for this quantity is unaffected by Q. □

Proof of Proposition 3. Appealing to Lemma 5, we establish the first point by analyzing b_L and b_R separately. First, when extended to the entire real line b_L and $b_{\text{BNE}}^{\text{UPA}}$ cross exactly twice. To see this, we define $\eta = (n-1)/(n-2)$ and check

$$\theta - \sqrt{(2\theta - \rho q)\,\rho q} = \theta - \eta\rho q \quad \Longleftrightarrow \quad \eta^2\rho^2 q^2 = (2\theta - \rho q)\,\rho q.$$

This equation has a trivial solution at $q = 0$, and a nontrivial solution at a unique $q_L^* > 0$. We note that for $q \in (0, q_L^*)$, $b_L(q) < b_{\text{BNE}}^{\text{UPA}}(q)$, while for $q > q_L^*$, $b_L(q) > b_{\text{BNE}}^{\text{UPA}}(q)$.

Second, when extended to the entire real line b_R and $b_{\text{BNE}}^{\text{UPA}}$ also cross exactly twice. To see this, we check

$$\frac{(Q-q)\,(2\theta - (q+Q)\,\rho)}{2Q} = \theta - \eta\rho q \quad \Longleftrightarrow \quad 2Q\eta\rho q = 2\theta q + \left(Q^2 - q^2\right)\rho.$$

This quadratic equation has solutions at $[(2\theta - 2\eta\rho Q) \pm \sqrt{(2\theta - 2\eta\rho Q)^2 + 4Q^2\rho^2}]/[2\rho]$. Comparing the discriminant to the leading term reveals that one solution is negative, hence there is a unique crossing point at a positive quantity, which we denote by $q_R^* > 0$. Lemma 5 then implies that for $q \in [0, q_R^*)$, $b_R(q) < b_{\text{BNE}}^{\text{UPA}}(q)$, and for $q > q_R^*$, $b_R(q) > b_{\text{BNE}}^{\text{UPA}}(q)$.

The leading result of Proposition 3 is a consequence of the following observation: either $b_{\text{MML}}^{\text{UPA}}(q) \le b_{\text{BNE}}^{\text{UPA}}(q)$ for all $q \in [0, Q]$, or there is a crossing point $\hat{q}$ at which $b_{\text{MML}}^{\text{UPA}}(\hat{q}) = b_{\text{BNE}}^{\text{UPA}}(\hat{q})$. The preceding arguments establish that $b_{\text{MML}}^{\text{UPA}}$ cannot cross $b_{\text{BNE}}^{\text{UPA}}$ from above for $q \in (0, \overline{Q})$, hence $b_{\text{MML}}^{\text{UPA}}(q) > b_{\text{BNE}}^{\text{UPA}}(q)$ for all $q \in (0, \overline{Q})$.

Now note that as n tends to infinity, $b_{\text{BNE}}^{\text{UPA}} \to \theta - \rho q = v(q;\theta)$ since $(n-1)/(n-2) \to 1$. Since bids under maximal uncertainty are always below value, BNE bids are higher. Similarly, when ρ is very close to 0 the BNE bid is essentially θ for all quantities. □

Proof of Proposition 4. BNE bids are arbitrarily close to θ for all q if ρ is sufficiently close to 0. These bids are therefore higher than bids under maximal uncertainty, which are boundedly far from θ when q is small, irrespective of ρ. Let $\overline{Q}/n < q^*$. Pycia and Woodward [2025, Theorem 1] implies that $b_{\text{BNE}}^{\text{PAB}}(\overline{Q}/n) = v(\overline{Q}/n;\theta)$. Hence, BNE bids for larger quantities are above value and therefore above the bids under maximal uncertainty. □

A.3.5 Comparison of equilibrium outcomes

Proof of Proposition 6. By Lemma 5, for all $q \in [0, \overline{Q}_i]$,

$$\theta - \sqrt{(2\theta - \rho q)\,\rho q} \geq \frac{(Q - q)\,(2\theta - (Q + q)\,\rho)}{2Q}.$$

Then we may bound revenue under maximal uncertainty by

$$\left(\theta - \sqrt{(2\theta - \rho q)\,\rho q}\right) nq.$$

Optimizing this expression with respect to q requires

$$\left(\theta - \sqrt{(2\theta - \rho q)\,\rho q}\right) n - \left[\frac{\theta - \rho q}{\sqrt{(2\theta - \rho q)\,\rho q}}\right] n\rho q = 0.$$

This expression may be rearranged to solve

$$(2\theta - \rho q)\,\theta^2 \rho q = (3\theta - 2\rho q)^2\,\rho^2 q^2.$$

Letting $x \equiv \rho q/\theta$, this is solved at

$$4x^3 - 12x^2 + 10x - 2 = 0. \tag{10}$$

We now consider the zeros of this cubic. Its derivative is $12x^2 - 24x + 10$, which has zeros at $1 \pm \frac{1}{12}\sqrt{24}$. Since Equation (10) is zero when $x = 1$, it follows that there is at most a single zero on the interval $x \in (0, 1)$; and, as the first-order condition for a profit-maximization problem, there will be exactly one solution in $(0, 1)$.

It suffices to show that (10) is positive at q such that $b^{\text{UPA}}_{\text{BNE}}(q) = b^{\text{UPA}}_{\text{MML}}(q)$. From the proof of Proposition 3, this quantity $q^\star$ is such that $2\theta = (1+\eta^2)\rho q^\star$. By the substitution $x = \rho q/\theta$ we check the sign of (10) at $x = 2/(1 + \eta^2)$. Since $\eta \in [1, 2]$ it is sufficient to check the sign of (10) for $x \in [2/5, 1]$. We check

$$4\left(\frac{2}{5}\right)^3 - 12\left(\frac{2}{5}\right)^2 + 10\left(\frac{2}{5}\right) - 2 = \frac{32}{125} - \frac{48}{25} + \frac{20}{5} - 2 = \frac{1}{125}\,(32 - 240 + 500 - 250) > 0.$$

Then (10) is positive at $x = 2/5$ and hence is positive for all $x \in [2/5, 1]$, by properties of the cubic established above. It follows that $q^\star_{\text{MML}}$ is to the left of q^*, and hence $b^{\text{UPA}}_{\text{MML}}(q^\star_{\text{MML}}) < b^{\text{UPA}}_{\text{BNE}}(q^\star_{\text{MML}})$. Then optimal revenue under maximal uncertainty is below optimal revenue in Bayes-Nash equilibrium. □

B Increasing marginal values

In this appendix we analyze bidding with increasing marginal values in the two-unit case. Marginal values are v_{i1} and v_{i2} and such that $0 \leq v_{i1} < v_{i2}$. Bidder i submits two bids b_{i1} and b_{i2}. By the auction rules, b_{i1} and b_{i2} are the expressed willingnesses to pay for the first and second units, respectively. Note that $b_{i2} > v_{i2}$ cannot be optimal; hence, let $b_{i2} \leq v_{i2}$.

B.1 Pay-as-bid auction

In the pay-as-bid auction, maximal loss is analogous to Equation (1):

$$\max\left\{\left((v_{i1} - b_{i1}) + (v_{i2} - b_{i1})\right)_+ ,\ (b_{i1} - b_{i2}) + (v_{i2} - b_{i2})\ ,\ b_{i1} + b_{i2}\right\};$$

the only difference is that when bidder i loses the auction but could have won at least one unit, winning two units is always better than winning just one. When marginal values are relatively similar and $b_{i1} > b_{i2}$, then the minimax-loss bid vector is as with decreasing marginal values: $b_{i1}^{\text{PAB}} = \frac{1}{9}\left(3v_{i1} + 2v_{i2}\right)$ and $b_{i2}^{\text{PAB}} = \frac{v_{i2}}{3}$.

Observe that this optimal bid vector is consistent with $b_{i1} \geq b_{i2}$ only if $3v_{i1} \geq v_{i2}$. In this case, the bid for the first unit is also less than v_{i1}. In the other case, the constraint $b_{i1} \geq b_{i2}$ is binding and the case in which nothing is won is not a worst case (because regret conditional on winning one unit is always higher).[34] The minimax-loss bid vector is

$$b_{i1}^{\text{PAB}} = \begin{cases} \frac{1}{9}\left(3v_{i1} + 2v_{i2}\right) & \text{if } 3v_{i1} \geq v_{i2} \\ \frac{v_{i2}}{3} & \text{if } 3v_{i1} < v_{i2} \end{cases} \text{ and } b_{i2}^{\text{PAB}} = \frac{v_{i2}}{3}.$$

Note that with strongly increasing marginal values, the bid is independent of v_{i1}.

B.2 First rejected bid uniform-price auction

In the first rejected bid uniform-price auction, the conditional regret minimizing bid $b_{i1}^{\text{FRB}} = v_{i1}$ and $b_{i2}^{\text{FRB}} = \frac{v_{i2}}{2}$ is optimal and feasible as long as $2v_{i1} \geq v_{i2}$. In general, maximal regret

[34] Because $b_{i1} = b_{i2} = b$ a bidder who wins a single unit can never simultaneously win a second unit and save payment for the first unit, and regret in this case is at least $v_{i2} - b$. If the bidder wins zero units, regret can only be maximized if she would prefer to win two units, in which case she would receive utility $v_{i1} - b < 0$ for the first unit and utility $v_{i2} - b$ for the second unit. Then receiving zero units always results in less loss than receiving a single unit.

equals

$$\max\left\{\underbrace{(v_{i1}-b_{i1}+v_{i2}-b_{i1})_+}_{\text{(i)}},\underbrace{(b_{i1}-v_{i1})_+,b_{i2},v_{i2}-b_{i2}}_{\text{(ii)}},\underbrace{2b_{i2}-v_{i1}-v_{i2},2b_{i2}-v_{i2}}_{\text{(iii)}}\right\}.$$

These are the regrets associated with (i) winning nothing but wanting to win two units; (ii) winning one unit but wanting to win zero, one, or two units, respectively; (iii) winning two units but wanting to win zero or one unit. Note that the last case dominates the second-last and that if $b_{i1}=b_{i2}$, then $b_{i1}=b_{i2}>b_{i1}-v_{i1}$. Moreover, if $b_{i1}=b_{i2}$ and $v_{i1}<b_{i1}$, then $v_{i2}-b_{i2}>v_{i1}-b_{i1}+v_{i2}-b_{i1}$. Hence, maximal loss is minimized by bidding

$$b_{i1}^{\text{FRB}}=\begin{cases}v_{i1} & \text{if } 2v_{i1}\geq v_{i2}\\ \frac{v_{i2}}{2} & \text{else}\end{cases}\quad\text{and } b_{i2}^{\text{FRB}}=\frac{v_{i2}}{2}.$$

As in the pay-as-bid auction, with strongly increasing marginal values minimax-loss bids are independent of v_{i1}.

B.3 Last accepted bid uniform-price auction

Recall that the conditional regret minimizing bid is $b_{i1}^{\text{LAB}}=(v_{i1}+v_{i2})/3$ and $b_{i2}^{\text{LAB}}=v_{i2}/3$ in the LAB uniform-price auction when the marginal values are sufficiently flat. Since $b_{i1}^{\text{LAB}}\geq b_{i2}^{\text{LAB}}$ for all values of v_{i1} and v_{i2}, the bid is optimal and feasible also for increasing marginal values.

C Bidpoint-constrained minimax-loss bids

The appendix contains the detailed analysis of the bidpoint-constrained case of Section 6 and discusses design implications.

C.1 Pay-as-bid auctions

Equating underbidding regret across all units leads to the following expression for minimax-loss bids.

Theorem 3 (Constrained minimax-loss bids in pay-as-bid)**.** *The unique minimax-loss bid* *in the constrained pay-as-bid auction solves*

$$\left(q^{PAB}, b^{PAB}\right) \in \underset{q', b'}{\operatorname{argmin}} \int_0^Q \left(v^i(x) - \hat{b}'(q_0)\right)_+ dx,$$
$$s.t. \int_0^{q'_k} \left(\hat{b}'(x) - \hat{b}'(q'_k)\right) dx + \int_{q'_k}^Q \left(v^i(x) - \hat{b}'(q'_k)\right)_+ dx$$
$$= \int_0^Q \left(v^i(x) - \hat{b}'(q_0)\right)_+ dx$$

for all $k = 1, 2, \ldots, M$.

Proof of Theorem 3. This proof is substantially similar to proof of the equivalent result for the unconstrained pay-as-bid auction (Lemma 4). As in the proof of Lemma 4, Lemma 2 implies that the loss minimization problem is

$$(q^\star, b^\star) \in \underset{(q', b')}{\operatorname{argmin}} \left[\max_{k \in \{0,1,\ldots,M\}} \left[\max \left\{ \overline{R}_{q'_k}\left(b'; v^i\right), \underline{R}_{q'_k}\left(b'; v^i\right) \right\} \right] \right].$$

By definition, $\underline{R}^{\text{PAB}}_{q_M}(b; v^i) \geq \overline{R}^{\text{PAB}}_{q_k}(b; v^i)$ for all k. Then the loss optimization problem in the pay-as-bid auction can be written

$$(q^\star, b^\star) \in \underset{(q', b')}{\operatorname{argmin}} \left[\max_{k \in \{0,1,\ldots,M\}} \underline{R}^{\text{PAB}}_{q_k}\left(b'; v^i\right) \right].$$

Recall that

$$\underline{R}^{\text{PAB}}_{q_k}\left(b'; v^i\right) = \int_0^{q_k} \left(\hat{b}'(x) - \hat{b}'(q_k)\right) dx + \int_{q_k}^Q \left(v^i(x) - \hat{b}'(q_k)\right)_+ dx.$$

Note that $\underline{R}^{\text{PAB}}_{q_k}$ decreases as q_k increases while, for all $k' > k$, $\underline{R}^{\text{PAB}}_{q_{k'}}$ increases as q_k increases. It follows that if $(q^\star, b^\star)$ is optimal, then $\underline{R}^{\text{PAB}}_{q_k}(b^\star; v^i) = \underline{R}^{\text{PAB}}_{q_{k'}}(b^\star; v^i)$ for all k, k'. □

We now show how to find the bidpoint-constrained minimax-loss bids in Example 4.

Calculations for Example 4. Equating conditional loss across units requires $R_{k+1} - R_k = 0$

for all k. This is

$$
\begin{aligned}
0 &= \left[\sum_{k'=0}^{k+1} (b_{k'} - b_{k+2})(q_{k'} - q_{k'-1}) + (Q - q_{k+1})(\theta - b_{k+2})\right] - \\
&\qquad \left[\sum_{k'=0}^{k} (b_{k'} - b_{k+1})(q_{k'} - q_{k'-1}) + (Q - q_k)(\theta - b_{k+1})\right] \\
&= (b_{k+1} - b_{k+2})(q_{k+1} - q_k) + (Q - q_{k+1})(\theta - b_{k+2}) \\
&\qquad + \sum_{k'=0}^{k} (b_{k+1} - b_{k+2})(q_{k'} - q_{k'-1}) - (Q - q_k)(\theta - b_{k+1}) \\
&= (b_{k+1} - b_{k+2}) q_{k+1} + (Q - q_{k+1})(\theta - b_{k+2}) - (Q - q_k)(\theta - b_{k+1}) \\
&= -Q b_{k+2} - (q_{k+1} - q_k)\theta + (Q + (q_{k+1} - q_k)) b_{k+1}.
\end{aligned}
$$

Let $g_k \equiv q_k - q_{k-1}$ be the gap between the k^{th} and $k+1^{\text{th}}$ bid points. Then we have

$$
\begin{aligned}
(Q + g_{k+1}) b_{k+1} = g_{k+1}\theta + Q b_{k+2} \quad &\iff \quad b_{k+1} = \frac{g_{k+1}}{Q + g_{k+1}}\theta + \frac{Q}{Q + g_{k+1}} b_{k+2} \\
&\iff \quad b_k = \frac{g_k}{Q + g_k}\theta + \frac{Q}{Q + g_k} b_{k+1}.
\end{aligned}
$$

We now solve recursively for optimal bids, conditional on bid points. When $k = M$, we have $b_{k+1} = 0$ by assumption, and $b_M = \frac{g_M}{Q+g_M}\theta$. For $k < M$, we have

$$
b_k = \sum_{k'=k}^{M} \frac{Q^{k'-k} g_{k'}}{\prod_{j=k}^{k'} [Q + g_j]} \theta.
$$

Since $R_0 = (\theta - b_1)Q$, the loss-minimization problem is (dropping the irrelevant constants θ and Q)

$$
\begin{aligned}
&\min_g 1 - \sum_{k=1}^{M} \frac{Q^{k-1} g_k}{\prod_{k'=1}^{k} [Q + g_{k'}]} \\
&= \min_g 1 - \frac{\sum_{k=1}^{M} \frac{1}{Q+g_k} \prod_{k'=k}^{M} [Q + g_{k'}] \, Q^{k-1} g_k}{\prod_{k'=1}^{M} [Q + g_{k'}]} \\
&= \min_g \frac{\prod_{k'=1}^{M} [Q + g_{k'}] - \sum_{k=1}^{M} \frac{1}{Q+g_k} \prod_{k'=k}^{M} [Q + g_{k'}] \, Q^{k-1} g_k}{\prod_{k'=1}^{M} [Q + g_{k'}]}.
\end{aligned}
$$

Denote the numerator by A_M. We show that $A_M = Q^M$. First, $A_1 = Q$:

$$A_1 = [Q + g_1] - \frac{1}{Q + g_1} [Q + g_1] g_1 = Q.$$

The result follows by induction on M; assuming $A_M = Q^M$, we have

$$\begin{aligned}
&\prod_{k'=1}^{M+1} [Q + g_{k'}] - \sum_{k=1}^{M+1} \frac{1}{Q + g_k} \prod_{k'=k}^{M+1} [Q + g_{k'}] Q^{k-1} g_k \\
&= [Q + g_{M+1}] \left[Q^M + \sum_{k=1}^{M} \frac{1}{Q + g_k} \prod_{k'=k}^{M} [Q + g_{k'}] Q^{k-1} g_k \right] - \sum_{k=1}^{M+1} \frac{1}{Q + g_k} \prod_{k'=k}^{M+1} [Q + g_{k'}] Q^{k-1} g_k \\
&= [Q + g_{M+1}] Q^M + \sum_{k=1}^{M} \frac{1}{Q + g_k} \prod_{k'=k}^{M+1} [Q + g_{k'}] Q^{k-1} g_k - \sum_{k=1}^{M+1} \frac{1}{Q + g_k} \prod_{k'=k}^{M+1} [Q + g_{k'}] Q^{k-1} g_k \\
&= [Q + g_{M+1}] Q^M - Q^M g_{M+1} = Q^{M+1}.
\end{aligned}$$

Then the loss minimization problem is

$$\min_g \frac{Q^M}{\prod_{k=1}^{K} [Q + g_k]}, \text{ s.t. } g_k \geq 0 \text{ and } \sum_{k=1}^{M} g_k \leq Q.$$

This is solved by $g_k = Q/M$. The resulting bids are

$$\begin{aligned}
b_{k|M} = \sum_{k'=k}^{M} \frac{Q^{k'-k} g_{k'}}{\prod_{j=k}^{k'} [Q + g_j]} \theta &= \sum_{k'=k}^{M} \frac{\frac{1}{M} Q^{k'-k+1}}{\prod_{j=k}^{k'} \left[\frac{M+1}{M} Q\right]} \theta \\
&= \sum_{k'=k}^{M} \frac{\frac{1}{M} Q^{k'-k+1}}{\left[\frac{M+1}{M} Q\right]^{k'-k+1}} \theta = \frac{\theta}{M} \sum_{k'=k}^{M} \left[\frac{M}{M+1}\right]^{k'-k+1}.
\end{aligned}$$

□

C.2 Uniform-price auctions

The following theorem proves that there is a unique minimax-loss bid in the constrained uniform-price auction.

Theorem 4 (Minimax-loss bids in constrained uniform-price auction)**.** *In the bidpoint-*

constrained uniform-price auction with M bid points, the unique minimax-loss bid solves

$$\left(q^{UPA}, b^{UPA}\right) \in \min_{q',b'} R,$$
$$s.t.\ q'_k b'_k = R \quad \forall k \in \{1, \ldots, M\},$$
$$and \int_{q'_{k-1}}^{Q} \left(v^i(x) - b'_k\right)_+ dx = R \quad \forall k \in \{1, \ldots, M\}.$$

Proof of Theorem 4. We first prove that the minimax bid (b_i, q_i) must solve

$$b_1 q_1 = b_k q_k \qquad \text{for } k \in \{1, 2, \ldots, M\}, \text{ and}$$
$$b_1 q_1 = \int_{q_{k-1}}^{Q} \left(v^i(x) - b_k\right)_+ dx \qquad \text{for } k \in \{1, 2, \ldots, M+1\}.$$

Let k denote the largest index for which maximal loss is attained, i.e, either $k = M+1$ if $\sup_{B^{-i}\in\mathcal{B}} L^{\text{UPA}}(b_i; B^{-i}, v^i) = \int_{q_M}^{Q} v^i(x)\, dx$ or

$$k = \max\left\{k' \colon \sup_{B^{-i}\in\mathcal{B}} L^{\text{UPA}}\left(b_i; B^{-i}, v^i\right) = \max\left\{\underline{R}^{\text{UPA}}_{q_{k'-1}}, \overline{R}^{\text{UPA}}_{q_{k'}}\right\}\right\}.$$

Let $k < M+1$. We show that $\underline{R}^{\text{UPA}}_{q_{k-1}} = \overline{R}^{\text{UPA}}_{q_k}$. Suppose $\underline{R}^{\text{UPA}}_{q_{k-1}} > \overline{R}^{\text{UPA}}_{q_k}$. As b_k appears in only these two expressions, raising b_k decreases only $\underline{R}_{q_{k-1}}$ and increases only $\overline{R}^{\text{UPA}}_{q_k}$. Suppose $\underline{R}^{\text{UPA}}_{q_{k-1}} < \overline{R}_{q_k}$. Decreasing b_k decreases $\overline{R}^{\text{UPA}}_{q_k}$ and increases $\underline{R}^{\text{UPA}}_{q_{k-1}}$. We do not have to worry about the effect on $\underline{R}^{\text{UPA}}_{q_k}$ as $\underline{R}^{\text{UPA}}_{q_k} < \overline{R}^{\text{UPA}}_{q_k}$.

Let $k = M+1$. Observe that $\int_{q_M}^{Q} v^i(x)\, dx = \underline{R}^{\text{UPA}}_{q_M} \leq \underline{R}^{\text{UPA}}_{q_{M-1}}$ as underbidding regret decreases in b_k and q_{k-1}. As regret is maximized by $M+1$, the inequality must hold with equality. The argument of the previous paragraph implies $\underline{R}^{\text{UPA}}_{q_{M-1}} = \overline{R}^{\text{UPA}}_{q_M}$. The result follows.

We now prove that a unique solution exists. To do so, note that we can express b_k as a function of q_{k-1} and q_k by solving

$$b_k q_k = \int_{q_{k-1}}^{Q} \left(v^i(x) - b_k\right)_+ dx$$

for b_k. The left-hand side increases in b_k and is 0 at $b_k = 0$. The right-hand side decreases in b_k, is positive for $b_k = 0$, and tends to 0 as b_k increases. Thus, there is a unique $b_k(q_{k-1}, q_k)$ that solves the equation. The bid $b_k(q_{k-1}, q_k)$ decreases in q_{k-1} and q_k.

We then proceed by expressing $q_{k'}$ as a function of q_1 by solving $b_1(q_0, q_1)q_1 = b_{k'}(q_{k'-1}, q_{k'})q_{k'}$ iteratively for $q_{k'}$, $k' \in \{2, 3, \ldots, M\}$. There is a unique $q_{k'}$ for each q_1. Finally, the condition $b_M(q_{M-1}(q_1), q_M(q_1))q_M(q_1) = \int_{q_M(q_1)}^{Q} v^i(x)\, dx$ pins down q_1. □

C.3 Design implications

Recall from the end of Section 6 that the minimax-loss bid in the constrained uniform-price auction drops to 0 at a quantity at which the minimax-loss bid in the constrained pay-as-bid auction is still positive under flat marginal values. The ambiguous revenue comparison is immediate.

Comparison 5 (Ambiguous revenue)**.** *Depending on the joint value distribution, both ex post and expected revenues can be higher in either constrained auction format.*

We illustrate the ambiguous revenue comparison in the following numerical example.

Example 6. We simulate bidpoint-constrained auction outcomes for different choices of the number of allowed bid points M. In the simulated auctions the available quantity is $Q = 100$, hence the locations of bidpoints correspond to percentage of aggregate supply. We vary the number of bidders from $n = 2$ to $n = 10$. Bidders' marginal values are flat, $v(q) = \theta$, where θ follows a truncated lognormal distribution with support $\theta \in [0.5, 2]$ and mean 1. For each number of allowed bid points, M, we first compute constrained minimax-loss bids in both the pay-as-bid and uniform-price auctions. In the pay-as-bid auction bids are obtained from the expressions in Example 4; in the uniform-price auction bids are obtained from the simple search procedure outlined in Section 6.2.

Figure 6 plots average auction revenue as a function of the number of bid points M. As expected, increasing the number of bidders increases the seller's expected revenue: the highest value of n independent draws increases in n in expectation. In general, revenue is ambiguous in the auction format and the number of bid points M. As observed in Examples 4 and 5, bidders in a pay-as-bid auction with a single bid point will bid half their value for the full market quantity, and bidders in a uniform-price auction with a single bid point will bid more than half their value for less than the full market quantity. Revenue in the pay-as-bid auction is therefore half the highest marginal value, while revenue in the uniform-price auction is more than half the second-highest marginal value. It follows that expected revenue will be higher in the pay-as-bid auction when both the number of bid points and the number of bidders are small.

Although average revenues may be ranked, reverse rankings can be observed ex post. Figure 6 also compares ex post revenues and depicts the share of simulated auctions in which uniform-price revenue is higher than pay-as-bid revenue. As the number of bidders increases, the share of auctions in which revenue is higher in the uniform-price auction increases. Low-revenue outcomes mainly appear in uniform-price auctions with two bidders, and these "collusive" outcomes are less likely when there are many bidders. The uniform-price auction dominates the pay-as-bid auction with ten bidders in terms of revenue in

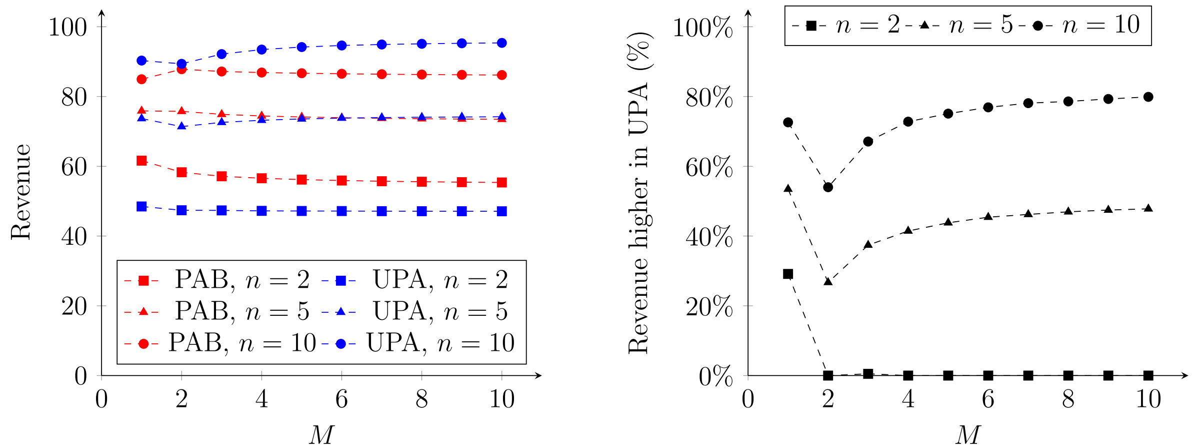

Figure 6: Average revenue (left) and ex post revenue comparison (right) as a function of number of bid points M.

expectation and ex post in the majority of auctions. The ambiguous, setting-dependent revenue ranking is in line with empirical results on multi-unit auctions.[35] Nonetheless, it is generally true that increasing the number of bidders increases the performance of the uniform-price auction relative to the pay-as-bid auction. Because initial bids are relatively high in the uniform-price auction and bids are relatively inelastic, increasing the number of bidders has strong upward influence on the market-clearing price, and thus on revenue. ◀

Figure 6 reveals that expected revenues can increase or decrease in the number of bid-points M. While a general analysis is beyond the scope of the paper, we provide the optimal M in two special cases.

Proposition 7. *When all bidders have flat marginal values, then the welfare-maximizing number of bid steps is* $M = 1$*. When there are infinitely many bidders, then revenue is maximized by* M *as large as possible.*

Proof. In the case of flat marginal values, it is efficient that the bidder with the highest type wins everything. Any auction selects the bidder with the highest type as the winner when $M = 1$.

When revenue is the objective and there are many bidders, then each bidder wins at most an arbitrarily small quantity. Since $b(0)$ then determines revenue and $b(0)$ is maximized by $M = \infty$ (because otherwise the bid is an average across lower values), the revenue-maximizing choice of M does not constrain the bidders. □

[35]See Pycia and Woodward [2025] for a summary of the ambiguous revenue rankings obtained in the empirical literature.

D Last accepted bid uniform-price auction

In this appendix, we provide the details on bidding in the last accepted bid uniform-price auction where bidders demand up to two units. Building on the analysis of the first rejected bid uniform-price auction in Section 3, we only provide the key steps of the analysis the uniform-price auction with the last accepted bid as the market-clearing price. We again restrict attention to bids below value: $b_{ij} \leq v_{ij}$ for $j = 1, 2$.

Case 1: zero units. As in the pay-as-bid auction, if the bidder wins zero units they know that they have underbid the two opponent bids. Their bids are most suboptimal if they could marginally increase their bid and win as many units as they desire, in which case loss is

$$\left[(v_{i1} - b_{i1}) + (v_{i2} - b_{i1})_+\right] - 0 = (v_{i1} - b_{i1}) + (v_{i2} - b_{i1})_+ .$$

The worst case bid distribution puts the two highest opponent bids marginally above b_{i1}.

Case 2: one unit. Conditional on winning one unit, the bidder overbids if b_{i1} sets the market-clearing price (and $c_1 = v_{i1}$ and $c_2 = 0$) and underbids if the market-clearing price is just above b_{i2} (and $c_1 = c_2 = v_{i2} + \epsilon$). In this case, loss is

$$\max\left\{b_{i1}, (v_{i2} - b_{i2})_+\right\} .$$

Case 3: two units. When the bidder wins two units, they set the market-clearing price. In this case, bids are most suboptimal when the bidder could have reduced bids to (almost) zero without losing any units; then loss is

$$[(v_{i1} - 0) + (v_{i2} - 0)] - [(v_{i1} - b_{i2}) + (v_{i2} - b_{i2})] = 2b_{i2}.$$

Maximal loss is then

$$\max\left\{(v_{i1} - b_{i1}) + (v_{i2} - b_{i1})_+ ,\ b_{i1},\ 2b_{i2},\ v_{i2} - b_{i2}\right\} .$$

Due to the different signs, maximal loss is minimized by equalizing at least *some* of the conditional losses; this contrasts the pay-as-bid auction, in which maximal loss is minimized by equalizing *all* of the conditional losses. Pairwise equalization of maximum loss gives a minimax-loss bid vector,

$$b_{i1}^{\text{LAB}} = \begin{cases} \frac{1}{3}\left(v_{i1} + v_{i2}\right) & \text{if } v_{i1} \leq 2v_{i2}, \\ \frac{1}{2}v_{i1} & \text{otherwise;} \end{cases} \quad \text{and } b_{i2}^{\text{LAB}} = \frac{v_{i2}}{3}.$$

The first bid can be found by equalizing the underbidding regret conditional on losing the auction $v_{i1} - b_{i1} + (v_{i2} - b_{i1})_+$ with the overbidding regret conditional on winning one unit b_{i1}. The second bid can be found by equalizing the underbidding regret conditional on winning one unit $v_{i2} - b_{i2}$ and the overbidding regret conditional on winning two units $2b_{i2}$.

While minimax-loss bids must minimize conditional regret for some unit, this will not in general determine the minimax-loss bid for all units. With demand for two units, worst-case loss minimization uniquely determines the bid for the first unit, but the bid for the second unit need only lie within the bounds $v_{i2} - L^{\text{LAB}} \leq b_{i2} \leq L^{\text{LAB}}/2$, where $L^{\text{LAB}} = b_{i1}$ is minimax loss in the uniform-price auction. Note the difference to the FRB auction where the last bid is uniquely determined by global regret minimization and the first bid from local regret minimization. As in the FRB uniform-price auction, we view the selection that optimizes the entire bidding function based on "local" worst cases natural.