# Self-Auditable Auctions

Kyle Woodward*

March 2020†

**Abstract**

We consider the amount of information necessary to verify that an auction has been run as claimed. A mechanism is *audited* by a post-allocation disclosure policy if each outcome maximizes the auctioneer's utility, conditional on consistency with the information released. One mechanism is more auditable than another if any disclosure policy that audits the latter also audits the former. When the seller cannot commit to any bounds on supply, only menus are auditable without additional disclosure. In contrast with other notions of auctioneer believability, fixed-supply discriminatory auctions are no more auditable than uniform price auctions. When supply is adjustable, the discriminatory auction is auditable without additional disclosure if the auctioneer claims to select an ex post profit-maximizing allocation, but the uniform price auction is not. Nonetheless, the ability to commit to a supply schedule via disclosure strictly improves auctioneer's expected revenue, even in the discriminatory auction.

## 1 Introduction

In multi-unit auctions, buyers may be skeptical of not only their payments but also their allocations. After bids are submitted, a seller may secretly adjust the quantity supplied, potentially improving his profits over the claimed mechanism.[1] Even if a bidder's payment is verifiable conditional on her allocated quantity, it is possible that the quantity she receives differs from that prescribed by the claimed mechanism.

In this paper, we consider what is necessary to make a bidder believe that an auction was run honestly, and how one auction format might be more inherently believable than another. Our analysis operates on post-auction information disclosure: after the auction is run, the auctioneer may publicly release information sufficient (or not) for the bidders to conclude that the mechanism was run as claimed. Information that is sufficient for this belief *audits* the auction. We show that

*UNC Chapel Hill; kyle.woodward@unc.edu. We would like to thank Gary Biglaiser, Justin Burkett, Uwe Dulleck, John Hatfield, Zhiguo He, Fei Li, Shengwu Li, George Mailath, Thayer Morrill, Jawwad Noor, Marek Pycia, Max Rosenthal, Larry Samuelson, and audiences at UNC Chapel Hill, QUT, INFORMS, and SEA for valuable discussions and feedback.

†Original version January 2019. Substantial revisions January 2020.

[1] For our initial results we constrain attention to claims of inelastic supply of a known quantity $Q$. "Mis-supply" also applies to the case of elastic supply, where the seller can allocate a quantity which differs from that prescribed by the elastic supply curve.

when no information is released after the auction is run, no common auction format is auditable, and all auditable mechanisms are menus. In many mechanisms the seller-optimal quantity depends on the submitted bids, and there is an incentive to deviate from claimed supply ex post. Some information is therefore necessary to audit standard auctions. Somewhat surprisingly, it is possible for a disclosure policy to audit a discriminatory auction but not a uniform price auction, and vice versa. We interpret this as implying that neither auction format is more auditable than the other.

Our initial analysis makes the assumption that the auctioneer may costlessly adjust supply after the auction is run. The assumption of costless supply is quite strong, and we later relax it to consider sellers with positive costs. When the seller claims to optimize supply ex post, the discriminatory auction is auditable without further disclosure, but the uniform price auction is not. However, even in this setting commitment has value: an auctioneer with a potentially unknown marginal cost curve can improve revenue with an appropriate disclosure policy. Thus, even though a claimed supply curve may be auditable without further disclosure, committing (via auditing disclosure) to a different supply curve can improve the seller's outcomes.

In our model, a disclosure policy is equivalent to a partition of feasible bid profiles. Although it is intuitive to consider what kind of information is released (for example, public statements about quantities allocated are sufficient to audit the discriminatory auction), it is mathematically more straightforward to work directly with what agents know about bid profiles.[2] Public information about actions is more specific than public information about outcomes: one outcome may arise from any of a set of action profiles, while one action profile yields a single outcome. Moreover, any mechanism can be audited by some public information about actions, but there are common mechanisms (VCG and Spanish auctions, for example) that cannot be audited by information about outcomes.[3] While we do not make any claim as to the mechanism for verifying public information, a natural interpretation is that if a market participant sees public information that conflicts with her own information, she reports the misinformation to an authority, and this is costly for the auctioneer.[4]

Our first contribution is a method by which the inherent believability of two mechanisms may be compared. A disclosure policy audits a mechanism if the mechanism's specified outcome maximizes the seller's utility within the set of all outcomes that are consistent with public and bidder-private information. One mechanism is *more auditable* than another if any disclosure policy that audits

[2]Practically, it may be natural to constrain disclosure policies to lie within a particular set, for example the set of information structures that identify an auction's market-clearing price. Restricting the space of permissible disclosure policies does not affect any intuitive arguments, but can affect the ranking of mechanism auditability. In Section 3.2 we consider the possibility that disclosure corresponds to a partition of outcomes.

[3]In Spanish auctions the transfer to the auctioneer is a personalized convex combination of bid discrimination and uniform pricing. Crucially, in both VCG and Spanish auctions a bidder's payment depends on non-marginal bids of her opponents.

[4]Under this interpretation it is natural to ask whether it is incentive compatible for a bidder to report an observed misallocation. When type spaces are sufficiently rich, there is no difference between exogenous reporting and incentive compatible reporting. If type spaces are relatively constrained, or if the auctioneer can use equilibrium information to determine whether to adjust a bidder's outcome, this equivalence breaks down. Roughly, tacit collusion between a bidder and the auctioneer is possible only if the auctioneer can conditionally deviate from the stated mechanism, using inference of bidder preferences to determine when it is appropriate to do so.

the latter also audits the former. We define auditability as a property of the full set of feasible bid profiles, not only those that arise in equilibrium. This is in contrast with the traditional analysis of seller incentives, which focuses on whether the seller wants to adhere to a claimed mechanism, given participant equilibrium strategies. We employ a stronger formulation for three reasons. First, equilibrium bids generally vary with the fundamentals of the model, and if only equilibrium-feasible bids are analyzed it is straightforward to construct a disclosure policy which audits one auction and not another. Since information policies are typically long-lived, given once and adhered to over multiple runs of the mechanism, auditability should hold for any action profile that may arise. Second, in multi-unit auctions the intersection of equilibrium bid sets across auction formats may be trivial [Ausubel et al., 2014, Pycia and Woodward, 2019, Burkett and Woodward, 2020b]: bids in discriminatory auctions are relatively flat while bids in uniform price auctions are relatively steep. This differs significantly from symmetric single-unit auctions, where the set of equilibrium first price bids is a subset of the set of equilibrium second price bids. When equilibrium bid sets are distinct, it is trivially possible to audit the outcomes generated by one mechanism's equilibrium bids while not auditing another mechanism's equilibrium outcomes. Then to provide a meaningful comparison of auditability across mechanisms, it is necessary to evaluate all feasible bid profiles. Finally, allowing for all feasible bid profiles is a plausible approximation when it is common knowledge that agents like quantity and money, but there is no additional information about the intensity of these preferences or a model of where bids come from.

Having defined auditability, we show that if a mechanism is auditable without public information about actions, it generates the same ex post revenue as some posted menu. If an agent's payment to the seller can depend on her opponents' actions and there is no information about her opponents' actions, the seller will pick an outcome that maximizes the agent's transfer. A mechanism is auditable in zero information if this manipulation is not to the seller's advantage. This will be the case when the agent's payment does not depend on her opponents' actions, and generates the same revenue as some mechanism in which the bidder's allocation is determined solely by her own action. This immediately implies that when supply is freely adjustable some public information is necessary to audit, e.g., a discriminatory auction, even in the single-unit case.[5]

Although auditability of a discriminatory auction requires some public information, it is natural to assume that discriminatory auctions are more auditable than uniform price auctions. The former can be audited with public information corresponding to quantity allocations, while the latter appear to require information about prices as well as quantities.[6] We show that this intuition does not survive the robust definition of auditability, and neither of the discriminatory or uniform price auctions is more auditable than the other. Without further constraints on public information, incomparability results from the seller's ability to commit to his own indifference across uniformly-

[5] This contrasts the case in which supply is exogenously determined: when the seller is constrained to sell a single unit, the outcome of a first price auction is inherently believable [Skreta, 2015, Akbarpour and Li, 2018].

[6] Once quantities are public information, each agent can verify that her discriminatory auction transfer is correct as it does not otherwise depend on her opponents' bids. In a uniform price auction the seller needs to further communicate that all agents are paying the same market-clearing price, and (potentially) that this price is correct.

priced allocations. A seller will always weakly prefer allocating greater quantities in a discriminatory auction (and will strictly prefer greater allocations when bids for supra-marginal units are strictly positive).[7] In a uniform price auction the seller can be indifferent between two quantity allocations, provided the larger allocation has a lower market-clearing price. As a result, uniform price auctions can be audited without full revelation of quantity allocations, and there is a disclosure policy that audits a uniform price auction but not the corresponding discriminatory auction.

Lastly, we consider the auditability of ex post revenue maximization within a claimed mechanism. If it is common knowledge that the seller will select a revenue-maximizing outcome after bids are submitted, a discriminatory auction with elastic supply can be audited without public information; this is not the case for uniform price auctions. This result further connects the results of [Skreta, 2015] and [Akbarpour and Li, 2018] to those in our main text: supply commitment is equivalent to perfect elasticity at the committed supply, while elastic supply given public knowledge of zero marginal costs is equivalent to allocating all units with positive demand, which is incompatible with an artificial supply constraint. We show that this intuition does not depend on common knowledge of the auctioneer's marginal cost curve. Nonetheless, committing (via a disclosure policy) to a particular supply curve can strictly improve the auctioneer's revenue in common auction formats. Revenue-maximizing supply curves typically differ from marginal cost curves, and without the ability to commit to an optimal supply curve the auctioneer has an ex post incentive to deviate from claimed supply; this potential deviation depresses bids and reduces revenue.

Our results present a new criterion for mechanism selection. If bidders must be assured that a mechanism is run according to specified rules, believability of outcomes will be a design constraint. Common multi-unit auction formats have ambiguous revenue and welfare comparisons [Ausubel et al., 2014, Pycia and Woodward, 2019, Burkett and Woodward, 2020b], and sellers with different preferences over revenue and surplus, or facing different bidder fundamentals, may implement different auction mechanisms. More trustworthy sellers will have greater latitude to select a preferred mechanism, and if believability is a design constraint, less-believable mechanisms will be implemented more frequently by less-corrupt entities. Thus auditability is relevant not only for a concerned designer, but also for empirical investigations into the selection of an auction format (see, e.g., [Brenner et al., 2009]).

## 1.1 Related literature

This paper sits at the intersection of three literatures: the analysis of multi-unit auctions, evaluating auctioneers' incentives "outside" of the auction mechanism, and the structure of mechanisms that minimize participant objections. Multi-unit auctions have been studied extensively, but little is

[7] In a discriminatory auction with weakly positive bids, the seller always weakly prefers to allocate more units than claimed. This relies on the assumption that the good in question is digital, and can be produced at zero marginal cost; the extensive literature on digital goods (see Goldberg et al. [2006], Bhattacharya et al. [2013], and others) examines the question of how to sell goods producible at zero marginal cost. Our results stand apart from the digital goods literature: we address the believability of a claim to utilize a particular mechanism, and do not consider the mechanism design question. In Section 4 we consider the possibility of nonzero marginal costs.

known about the computation of equilibrium strategy profiles [Hortaçsu and Kastl, 2012]. In certain cases equilibrium expected revenues are directly computable [Back and Zender, 1993, Engelbrecht-Wiggans and Kahn, 1998, 2002, Wang and Zender, 2002, Holmberg, 2009, Burkett and Woodward, 2020b, Pycia and Woodward, 2019], but in general there is no theoretical comparison of auctioneer outcomes across different multi-unit auction formats [Ausubel et al., 2014]. This ambiguity is not clearly resolved under empirical investigation [Armantier and Sbaï, 2006, Castellanos and Oviedo, 2008, Kang and Puller, 2008, Armantier and Sbaï, 2009, Hortacsu and McAdams, 2010, Hortaçsu et al., 2018]. A reasonable takeaway is that there is no simple model of bidder values that provides an identical revenue comparison across observed implementations. We contribute to this literature by arguing that auditability may be a design constraint, affecting the ability of the auctioneer to implement a revenue- or efficiency-maximizing auction format. In some circumstances, this may constrain practical mechanism selection. The multi-unit auction literature implies that, when bidders demand more than one unit, the structure of observed bids varies strongly with auction format [Ausubel et al., 2014, Burkett and Woodward, 2020b, Pycia and Woodward, 2019]. This extends beyond the single-unit intuition that bids are higher in a second price auction than in a first price auction, since it is known that bid *curves* are less elastic in discriminatory auctions than in uniform price auctions; in single-unit auctions equilibrium bid spaces are nested, while in multi-unit auctions this is not the case. Incomparable equilibrium bid spaces inform our construction of disclosure policies, and the idea that auctions must be auditable for all bid profiles, not just those which arise in equilibrium.

The literature on adjustable supply in multi-unit auctions considers the possibility that the seller does not commit to a quantity to sell. In single-unit auctions, revenue can be improved with the introduction of a reserve price. In multi-unit auctions, this generalizes to a supply schedule. In uniform price auctions an elastic supply curve can improve revenue by inducing competition at low prices [LiCalzi and Pavan, 2005, Burkett and Woodward, 2020a]. Committing to a supply curve which increases in market-clearing price shrinks the set of equilibrium market-clearing prices, and can eliminate severe equilibrium underpricing. A similar result holds when an auctioneer selects aggregate quantity after bids are submitted [McAdams, 2007]. We consider the possibility of both elastic and adjustable supply, and show that the believability ex post profit maximization depends on the disclosure policy.

Our results are of a piece with analyses of a designer's incentives to report truthfully [Bester and Strausz, 2000, 2001, Akbarpour and Li, 2018]. Ensuring the auctioneer reports truthfully (in our case, selects the outcome dictated by the mechanism) is an equilibrium concept given bidder strategies. [Bester and Strausz, 2000] and [Bester and Strausz, 2001] focus on applicability of the revelation principle in the face of designer incentives. In these models designer utility depends on the profile of agent types, so designer incentive compatibility relies on inference from observed play. We focus on inference from disclosed information independent of bidder equilibrium, so there is not a direct tie to these results. Credibility [Akbarpour and Li, 2018] sharpens this notion, where a mechanism is credible if the auctioneer cannot improve his own outcome without alerting the agents. In some cases our notion of auditability provides a hierarchy of mechanisms, augmenting the binary concept of credibility; the previous note on equilibrium inference also applies.[8]

Noncommitment in auctions has been analyzed extensively [McAfee and Vincent, 1997, Skreta, 2006, McAdams and Schwarz, 2007b, Vartiainen, 2013, Skreta, 2015]. Broadly, these papers ask how commitment benefits a seller; alternatively, how can a seller improve his outcome by using information revealed by credulous bidders? [Calzolari and Pavan, 2006] and [Skreta, 2011] address a similar question, assuming commitment: how can value-relevant information disclosure within a game can affect outcomes and reporting incentives? Here we consider only disclosure after the auction is completed, and our auctioneer has a comparatively limited ability to deviate from the specified mechansim: can he improve his outcome without tipping off bidders to his deviation, independent of the information he learns from equilibrium bidding?[9]

Finally, our definition of auditability ties to work on finding mechanisms that minimize (valid) participant complaints. [Abdulkadiroğlu and Sönmez, 2003] demonstrates that stable matchings eliminate justified envy, in the sense that any outcome an agent prefers is infeasible according to given priorities and preferences. Given the trade-off between efficiency and equity, it may be possible to obtain more efficient outcomes while allowing for justified envy that cannot be acted upon. [Cantala and Pápai, 2014] and [Alcalde and Romero-Medina, 2015] study the related concepts of reasonable stability and $\tau$-fairness, relating to the ability of rematches implied by justified envy to themselves generate justified envy. [Troyan et al., 2018] requires that a proposed rematching not initiate a rejection chain that invalidates the rematching. [Ehlers and Morrill, 2018] terms a school matching legal if any student with justified envy cannot receive a better outcome in any other legal matching. As in our paper, this literature takes as given that complaints are to be avoided and remains generally agnostic as to why this is.

We continue in Section 2 by defining our model. Section 3 provides results, and shows that common multi-unit auction formats cannot be ranked in auditability without additional assumptions. Section 4 considers two extensions, in which the auctioneer may openly adjust supply, or may have a nontrivial marginal cost curve. Proofs not contained in the main text are given in the appendix.

[8] An extensive literature has examined the ability to audit claims ex post [Townsend, 1979, Mookherjee and Png, 1989, Mylovanov and Zapechelnyuk, 2017, and others], and mechanism design with evidence [Postlewaite and Schmeidler, 1986, Bull and Watson, 2007, Deneckere and Severinov, 2008, Ben-Porath and Lipman, 2012, and others]. In this paper we ask how to issue evidence to render a mechanism believable, rather than how to implement a mechanism subject to the evidence available, and assume that auditing occurs costlessly after the mechanism is run.

[9] [McAdams and Schwarz, 2007b] shows that the cost of non-believability is borne entirely by the seller; this result is echoed by [McAdams and Schwarz, 2007a] as well as our Proposition 3. A key distinction is that [McAdams and Schwarz, 2007b] and [McAdams and Schwarz, 2007a] explicitly model costs, while in our model all costs of misallocation are endogenous. This suggests that auctioneers have a vested interest in their mechanisms being auditable, and if information release is generally undesirable, more (inherently) auditable mechanisms will be preferred to less auditable mechanisms.

## 2 Model

We model a static multi-unit auction for quantity $Q \in \mathbb{N}$. There are $n$ bidders, $i \in \{1, \ldots, n\}$, and each bidder has private type $\theta_i \in \Theta_i$; the type profile $\theta = (\theta_i)_{i=1}^n \sim F$. An outcome $o'$ is an $n$-tuple of quantity-transfer pairs, $o' \in (\mathbb{N} \times \mathbb{R})^n \equiv O$, where $O$ is the feasible outcome space. We interpret this as the seller being able to implement arbitrary transfers, and supply arbitrary whole quantities.

After observing her type $\theta_i$, bidder $i$ submits bid $b_i \in B$. The aggregate bid profile is $b = (b_i, b_{-i})$. A bid profile is associated with a personal outcome $o^i(b) = (q^i(b), t^i(b))$,[10] and the outcome profile $o(b) = (o^1(b), \ldots, o^n(b)) \in O$ determines quantity allocations and transfers for each bidder. Because we hold the set of feasible actions fixed across auction formats, we refer to $o$ as a *mechanism*.[11] Bidder $i$'s realized utility is

$$u^i\left(b_i, b_{-i}; \theta_i\right) = \bar{u}^i\left(q^i\left(b_i, b_{-i}\right), t^i\left(b_i, b_{-i}\right); \theta_i\right).$$

We assume that $\bar{u}^i$ is weakly increasing in $q$ and strictly decreasing in $t$. Denote the auctioneer's utility by $u^0 : B^n \to \mathbb{R}$. We assume that the auctioneer cares only about revenue, so $u^0(b) = \sum_{i=1}^n t^i(b_i, b_{-i})$.

Prior to the auction, the auctioneer announces a *disclosure policy* $\mathcal{I} : B^n \rightrightarrows B^n$. The disclosure policy is a claim regarding which public disclosure will be made, conditional on a submitted bid profile. The auctioneer cannot commit to a disclosure policy, and its announcement represents only the manner in which post-allocation information is interpreted.[12] Announced disclosure must be consistent, $b \in \mathcal{I}(b)$, and must yield a partition of $B^n$, so that for any $b$ and $b'$, either $\mathcal{I}(b) = \mathcal{I}(b')$ or $\mathcal{I}(b) \cap \mathcal{I}(b') = \emptyset$. Bidder $i$'s own bid and public information yield the set of explicable outcomes.

**Definition 1. [Explicable outcome]** Outcome $o'_i$ is *explicable for bidder $i$*, given information $I$, if there is some $b'_{-i} \in B^{n-1}$ such that $o'_i = o^i(b_i, b'_{-i})$, and $(b_i, b'_{-i}) \in I$. Outcome profile $o'$ is *explicable*, given information $I$, if $o'_i$ is explicable for each bidder $i$.

Denote the set of outcomes explicable for bidder $i$, given information $I$, by $X^i(b; I)$, and the set of explicable outcomes, given information $I$, by $X(b; I) = \times_{i=1}^n X^i(b_i; I)$.[13] Note that if information $I$ is inconsistent with some agent's bid, $X(b; I)$ will be empty. We therefore say that $I$ is *consistent* if $X(b; I) \neq \emptyset$. If public information is truthful, $I = \mathcal{I}(b) \ni b$ and $o(b)$ is explicable for each bidder $i$, so $X(b; I) \supseteq \{o(b)\}$ is nonempty. Importantly, the set of explicable outcomes $X$ depends not only on the public information $I$ but also on the claimed mechanism $o$; holding fixed public

[10] We denote functions with superscripts and values with subscripts. For example, $t_i$ is the transfer paid by bidder $i$, and $t^i : B^n \to \mathbb{R}_+$ is the function mapping submitted bids to bidder $i$'s transfer.

[11] With the set of feasible actions fixed, the only difference between two mechanisms is the mapping from actions (messages) to outcomes. Then in our model, a mechanism $(B, o)$ is completely determined by $o$.

[12] Most of our results are unaffected by lack of commitment to an information policy. The seller's incentive to mis-allocate is similar to her incentive to misrepresent information, and the question of "how much" information is needed to believe a mechanism differs only slightly between the commitment and no-commitment cases.

[13] An outcome $o'$ is explicable given information $I$ if it is explicable for each bidder, but it may not be the case that there is $b' \in I$ such that $o' = o(b')$.

information $I$, the set of explicable outcomes will in general differ from mechanism to mechanism. Finally, define the set of explicable outcomes (independent of the information released) by $\mathcal{X}(b) = \cup_{I \in \text{Range}\,\mathcal{I}} X(b; I)$.

Explicability determines the set of deviations from plan the seller could implement, conditional on a particular bid profile. This, in turn, defines the auctioneer's ability to respond to incentives.

**Definition 2.** **[Audited mechanism]** The mechanism $o$ is *audited by* disclosure policy $\mathcal{I}$ if for all $b \in B^n$,

$$o\left(b\right) \in \underset{o' \in \mathcal{X}(b)}{\operatorname{argmax}}\, u^0\left(o'\right).$$[14]

That is, the auction is audited by disclosure policy $\mathcal{I}$ if the outcome generated by the auction maximizes the auctioneer's utility, from the set of all outcomes which are consistent with some public information $I$ and each bidder's knowledge of her own action $b_i$. Each bidder has reason to believe that such outcomes are honestly determined, as anything the seller could do to improve his own utility would lie outside the set of explicable outcomes.

*Remark* 1. Any mechanism $o$ is audited by the policy of full disclosure, $\mathcal{I}(b) = \{b\}$. All public information sets $I \in \text{Range}\,\mathcal{I}$ are atomic, $I = \{b'\}$. Then if $I \neq \mathcal{I}(b) = \{b\}$, some agent $i$ sees $b'_i \neq b_i$, and $X^i(b; I) = \emptyset$. Then $X(b; I) \neq \emptyset$ only if $I = \mathcal{I}(b)$, and $\mathcal{X}(b) = \{b\}$. Then $o(b)$ is the unique explicable outcome given full disclosure, and the mechanism is audited by $\mathcal{I}$.

The "for all" public information release requirement of auditability implies that if $\mathcal{I}$ audits $o$, any refinement of $\mathcal{I}$ audits $o$.

*Remark* 2. Suppose that $\mathcal{I}$ audits $o$, and that $\mathcal{I}'$ is such that $\mathcal{I}'(b) \subseteq \mathcal{I}(b)$ for all bid profiles $b \in B^n$. Then $\mathcal{I}'$ audits $o$. Otherwise, there is some $b \in B^n$ such that, under disclosure policy $\mathcal{I}'$, the seller prefers to implement $o' \neq o(b)$; this is supported by some public information $I' \in \text{Range}\,\mathcal{I}'$. The allocation $o'$ is also explicable (given information $I = \mathcal{I}(I')$) under disclosure policy $\mathcal{I}$, hence the seller prefers to implement $o' \neq o(b)$, contradicting auditability.

The auction game being audited by a disclosure policy is distinct from credibility of the meta-game in which the auction is augmented by disclosure. Credibility (in the sense of [Akbarpour and Li, 2018]) presumes that the auctioneer is responding to bidder strategies; in our case, auditing an auction is a feature of *all* feasible bid profiles, not just those which may arise in equilibrium. In this sense auditability is a stronger notion than credibility. However, our analysis is complementary: in what follows we focus on the properties of information necessary to yield a believable mechanism, rather than innate features of the mechanism per se. As noted in [Burkett and Woodward, 2020b] and our introduction, equilibrium bids take very different forms in different multi-unit auction formats, so requiring auditability of only equilibrium-feasible bids would void the ability to compare the auditability of different auction formats.

[14]For notational simplicity, we assume that $\operatorname{argmax}_{x \in \emptyset} f(x) = -\infty$, so that the seller's announcement and allocation must be minimally consistent.

Although disclosure policies are mathematically structured as partitions on the space of bid profiles, in practice information is related to standard market outcomes. For example, the market-clearing price is typically announced after the auction closes. We view this announcement as equivalent to announcing that the submitted bid profile was in the set of bid profiles that generated this market-clearing price. Additionally, there are (deterministic) mechanisms that cannot be audited with information about outcomes, even allowing for inference from one's own submitted bid.[15] Allowing for nondeterministic mechanisms, explicit information regarding the process generating random outcomes may be necessary; we account for some sources of randomness in Section 4.1.

Finally, the ability of a disclosure policy $\mathcal{I}$ to audit a mechanism $o$ offers a natural comparison of the auditability of two mechanisms. Whichever requires *more* disclosure, or a stronger commitment, is *less* auditable.

**Definition 3.** [**More auditable**] Auction $o$ is *more auditable* than auction $o'$ if for any disclosure policy $\mathcal{I}$ that audits $o'$, $\mathcal{I}$ audits $o$, and we write $o \trianglerighteq o'$. If $o \trianglerighteq o'$ and there is an disclosure policy $\mathcal{I}$ that audits $o$ and not $o'$, then auction $o$ is *strictly more auditable* than auction $o'$, and we write $o \triangleright o'$. If $o \ntrianglerighteq o'$ and $o' \ntrianglerighteq o$, $o$ and $o'$ are *audit-incomparable*.

**Definition 4.** [**Zero auditability**] The mechanism $o$ is *zero-auditable*, or *auditable in zero information*, if $\mathcal{I} \equiv B^n$ audits $o$.

# 3 Results

We now use our auditability ordering to show that menu mechanisms are more auditable than any other mechanism, and are the revenue-maximizing mechanisms auditable in zero information.

**Definition 5.** [**Menu mechanisms**] A mechanism $o$ is a *menu mechanism* if bidder $i$'s outcome depends only on her own action, $o^i(b_i, b_{-i}) \equiv o^i(b_i)$.

It is straightforward to see that traditional posted-price mechanisms are menu mechanisms. The more general form of a menu mechanism allows for quantity-dependent pricing. Because bidder $i$'s outcome depends only on her own bid, it is immediate that menu mechanisms are audited by any information policy $\mathcal{I}$.

**Theorem 1.** [**Menus are zero-auditable**] *Let $\mathcal{I}$ be any information policy. If $o$ is a menu mechanism, then $\mathcal{I}$ audits $o$.*

*Proof.* For any bidder $i$, bid $b_i$, and outcome $o'$, $o'_i$ is explicable for bidder $i$ if there is $b'_{-i}$ such that $(b_i, b'_{-i}) \in I$ and $o'_i = o^i(b_i, b'_{-i})$. Since $o^i(b_i, b'_{-i}) \equiv o^i(b_i)$, $o'_i$ is explicable if and only if $o'_i = o^i(b_i)$. Then given any public information $I$, the set of explicable outcomes $X(b; I)$ has $o'_i = o^i(b)$ (or is empty) for all bidders $i$, and hence $X(b; I) \in \{\{o(b)\}, \emptyset\}$. Then $o(b) \in \arg\max_{o' \in \{o(b)\}} u^0(o')$. $\square$

[15] In Spanish treasury auctions transfers are computed as a combination of discriminatory and uniform price, where bids above the average winning bid pay the average winning bid and bids below the average bid are discriminated against. Verifying that the appropriate transfer was made requires verifying the average winning bid, which requires substantial information about opponent bid curves, and outside of simple cases cannot be derived from outcomes and own bid.

**Corollary 1. [Audit-dominance of menus]** *If $o$ is a menu mechanism, then $o \trianglerighteq o'$ for any mechanism $o'$.*

To enable comparison across mechanisms our auction model assumes identical bid spaces, independent of the outcome function. We therefore have no innate ability to distinguish between, e.g., a posted price mechanism and a discriminatory auction in which bids are always flat at a constant price, up to some endogenous quantity. Rather than make claims of uniqueness, we show that any mechanism auditable in zero information is outcome-equivalent to some menu mechanism.

**Theorem 2. [Zero-auditable mechanisms are menus]** *Let $o$ be a mechanism. If $o$ is audited by all disclosure policies $\mathcal{I}$, then $o$ is revenue-equivalent to a menu mechanism.*

*Proof.* Let $\mathcal{I}(b) = B^n$. Then any outcome $o'$ is explicable so long as it is consistent with each participant's bid, $o'_i \in o^i(b_i, B^{n-1})$ for each bidder $i$. Recall that $u^0$ is linear in transfers, so

$$o\left(b\right) \in \underset{o' \in \mathcal{X}(b)}{\operatorname{argmax}}\, u^0\left(o'\right) = \underset{\forall i,\ o'_i \in o^i(b_i, B^{n-1})}{\operatorname{argmax}} \sum_{i=1}^{n} t'_i.$$

Public information $I$ is irrelevant, since $\mathcal{I}(b) = B^n$ for all $b$. Maximizing $t'_i$ depends only on the constraint imposed by $b_i$, so $\max_{\forall i,\ o'_i \in o^i(b_i, B^{n-1})} \sum_{i=1}^{n} t'_i = \sum_{i=1}^{n} \max_{o'_i \in o^i(b_i, B^{n-1})} t'_i$. Then since $\mathcal{I}$ audits $o$, $t^i(b_i, b_{-i})$ is independent of $b_{-i}$. $\square$

**Corollary 2. [Menus are revenue-maximizing]** *Menu mechanisms are revenue-maximizing in the class of zero-auditable mechanisms.*

Theorem 2 distinguishes credibility from auditability in our model. Under known supply the first price auction is credible [Akbarpour and Li, 2018], however in our model with adjustable supply the first price auction is not auditable in zero information. In many multi-unit auction contexts there is no natural cap on the supply available for auction; or, if such a cap exists, it is large enough that bids will be nonaggressive.[16] Consider a seller who claims to be selling a single unit in a first-price auction. If bidders are mistaken in their belief that there is a single unit, the seller can solicit relatively aggressive bids, then award each bidder a unit, receiving higher revenue than if he had abided his quantity commitment and sold only a single unit.[17]

### 3.1 Common multi-unit auctions

We now define three common multi-unit auction formats. In each case, we assume the auctioneer is claiming to sell $Q \in \mathbb{N}$ units. The auctioneer solicits weakly decreasing, positive bid vectors

[16] For example, in a first price auction between two bidders, each of whom demands one unit, equilibrium bids are zero when two units are supplied.

[17] [Akbarpour and Li, 2018] shows that this format can be credible if the seller is truly and publicly quantity-constrained, or (equivalently) faces infinite marginal costs above a particular quantity. For a translation to auditability, see Corollary 3. Exogenously constraining quantity means that the auctioneer's maximization problem cannot be separated, as in the proof of Theorem 2, since allocating a unit to bidder $i$ means not allocating the same unit to bidder $j$. Because we relax this constraint, all zero-information auditable mechanisms are menus.

$b \in \mathbb{R}_{+}^{Q}$, and awards units to the $Q$ highest bids.[18] When tiebreaking is necessary, we assume a deterministic tiebreaking rule is used. Given supply $Q$ and bid profile $b$, the last accepted bid and first rejected bid are

$$b_{\mathrm{LA}} = \inf \{p : \# \{(i, q) : b_{iq} \geq p\} < Q\}, \quad b_{\mathrm{FR}} = \inf \{p : \# \{(i, q) : b_{iq} > p\} \leq Q\}.$$

The last accepted bid is the lowest price at which the market has weak excess supply, and the first rejected bid is the highest price at which the market has strict excess demand.[19] Accordingly, we refer to the highest and lowest market-clearing prices $p^{\star}_{\mathrm{LAB}}$ and $p^{\star}_{\mathrm{FRB}}$, and a price $p^{\star}$ is a *market-clearing price* if $p^{\star}_{\mathrm{LAB}} \geq p^{\star} \geq p^{\star}_{\mathrm{FRB}}$. In standard multi-unit auctions all bids $b_{iq} > p^{\star}_{\mathrm{FRB}}$ are awarded, and tiebreaking is necessary only if $p^{\star}_{\mathrm{LAB}} = p^{\star}_{\mathrm{FRB}}$. Denote by $q^{i}(b)$ the quantity obtained by bidder $i$ when the aggregate bid profile is $b \in B^{n}$.

In the discriminatory (or pay-as-bid) auction, bidders pay their bids up to their quantity allocation,

$$o^{i}_{\mathrm{PAB}}(b) = \left(q^{i}(b), t^{i}_{\mathrm{PAB}}(b)\right), \quad t^{i}_{\mathrm{PAB}}(b) = \sum_{q=1}^{q^{i}(b)} b_{iq}.$$

In a uniform price auction, bidders pay a market-clearing price for each unit they receive. Following the analysis in [Burkett and Woodward, 2020b] we consider the last accepted bid and first rejected bid uniform price auctions, where

$$\begin{aligned} o^{i}_{\mathrm{LAB}}(b) &= \left(q^{i}(b), t^{i}_{\mathrm{LAB}}(b)\right), & t^{i}_{\mathrm{LAB}}(b) &= p^{\star}_{\mathrm{LAB}}(b)\, q^{i}(b); \\ o^{i}_{\mathrm{FRB}}(b) &= \left(q^{i}(b), t^{i}_{\mathrm{FRB}}(b)\right), & t^{i}_{\mathrm{FRB}}(b) &= p^{\star}_{\mathrm{FRB}}(b)\, q^{i}(b). \end{aligned}$$

Auditability of an auction format depends on the source of the seller's misallocation incentive. In a discriminatory auction it must be clear that no more than the claimed quantity $Q$ may be profitably sold. Because the seller's profits are weakly increasing in quantity allocated, it is never necessary to assert that no less than $Q$ was sold: if selling $Q' < Q$ is profit-maximizing given bid profile $b$, selling $Q$ generates the same revenue and therefore $o(b) \in \mathrm{argmax}_{o' \in \mathcal{X}(b)}\, u^{0}(o')$. Then if $\mathcal{I}$ audits $o_{\mathrm{PAB}}$, $\mathcal{I}$ need only make explicit aggregate allocations, and then only in certain circumstances. Announcing quantity allocations is sufficient to audit the discriminatory auction.[20]

**Lemma 1. [Quantities audit PAB]** *Suppose that $\mathcal{I}$ is such that $q(b') = q(b)$ for all $b' \in \mathcal{I}(b)$. Then $\mathcal{I}$ audits $o_{PAB}$.*

[18] Because auditability of a mechanism must hold for any feasible bid profile, in some of our arguments (but not the results) we constrain attention to the case of $n = 2$ bidders. If one common multi-unit auction format does not audit-dominate another when there are only two bidders, it follows that it does not audit-dominate the other when there are more bidders. Bidders $i \in \{1, 2\}$ can submit the bid profiles implying audit non-dominance, and bidders $i \in \{3, 4, \ldots, n\}$ can bid zero and remain out of the auction. Constraining attention to the case of $n = 2$ bidders is therefore without loss of generality in many cases.

[19] The mathematical expressions deviate from this intuition to handle the possibility of tiebreaking.

[20] It is sufficient but not necessary to exactly announce quantity allocations to audit a discriminatory auction. For example, when bids are low enough that the market-clearing price is zero, it is sufficient to announce that the market-clearing price is zero and that quantities do not exceed a given amount.

*Proof.* Suppose that $o'_i$ is explicable for agent $i$, given bid profile $b$ and information $I = \mathcal{I}(b)$. Since $q(b') = q(b)$ for all $b' \in \mathcal{I}(b)$, $q'_i = q^i(b)$. In the discriminatory auction, bidder $i$'s transfer is independent of her opponents' bids, hence $t'_i = \sum_{q=1}^{q'_i} b_{iq} = t^i_{\text{PAB}}(b)$. Then $o'_i = o^i(b)$ and there is a unique explicable outcome conditional on information $I$, $X(b; I) = \times_{i=1}^{n} \{o^i(b)\} = \{o(b)\}$. Furthermore, by assumption all partitions $I \in \text{Range}\,\mathcal{I}$ are such that $\sum_{i=1}^{n} q^i(b') = Q$ for all $b' \in I$. Misrepresenting bid information can only reallocate quantity, not alter the aggregate quantity allocated; since transfers are discriminatory, it is weakly optimal for the auctioneer to disclose the correct partition $\mathcal{I}(b)$. Then $o(b) \in \text{argmax}_{o' \in \mathcal{X}(b)} u^0(o')$, and $\mathcal{I}$ audits $o_{\text{PAB}}$. $\square$

There is a natural sense in which it may seem that less information is required to verify a discriminatory auction than a uniform price auction. In a discriminatory auction the seller's own incentives ensure that knowledge of quantity is sufficient to know that outcomes were honestly computed. In a uniform price auction knowledge of quantity remains critical, but it is also essential to verify that each agent is paying the same market-clearing price. An agent asked to make a transfer equal to their last accepted bid does not know if this bid is the *aggregate* last accepted bid, but observing that each opponent is paying the same per-unit price is sufficient to confirm this outcome.

Surprisingly (in light of this argument) the last accepted bid auction is not less auditable than the discriminatory auction. This follows from the standard monopolist's problem: a discriminatory monopolist facing zero marginal costs wants to maximize quantity sold, while a posted-price monopolist may have multiple optimal prices.

**Theorem 3. [PAB audit-incomparable with LAB]** *When $Q > 1$, the discriminatory auction is no more auditable than the last accepted bid uniform price auction and vice versa, $o_{PAB} \ntrianglerighteq o_{LAB}$ and $o_{LAB} \ntrianglerighteq o_{PAB}$.*

Proving Theorem 3 requires showing that a disclosure policy can audit a discriminatory auction but not a last accepted bid auction, and (different) information can audit a last accepted bid auction but not a discriminatory auction. The former direction is straightforward. Consider a disclosure policy that announces only the aggregate allocation $q$. By Lemma 1 this disclosure policy audits the discriminatory auction. However, this disclosure policy leaves the market-clearing price $\overline{p}^\star$ unspecified. For any given bidder, observing a per-unit price equal to their marginal bid $b^i(q_i)$ is consistent with their allocation and the last accepted bid pricing rule. In general bidders will have different marginal bids, and the auctioneer can explicably charge bidders different per-unit prices.

Showing that a disclosure policy can audit a last accepted bid auction but not a discriminatory auction is simplest in a parameterized example. Suppose that the auctioneer claims that $Q = 2$ units are available, and $n = 2$ bidders are participating in an auction. After bids are submitted the auctioneer discloses the full profile of bids, unless bidder 1 bids $b_1 = (2, 1)$ and bidder 2 bids $b_2 = (2, 0)$ or $b_2 = (0, 0)$. Under this disclosure policy the auctioneer can manipulate outcomes only if $b_1 = (2, 1)$ and $b_2 \in \{(2, 0), (0, 0)\}$. If $(b_1, b_2) = ((2, 1), (2, 0))$ bidder 1 obtains $q_1 = 1$ at a clearing price of $\overline{p}^\star = 2$, while if $(b_1, b_2) = ((2, 1), (0, 0))$ bidder 1 obtains $q_1 = 2$ at a clearing price of $\overline{p}^\star = 1$.

In either case bidder 1's last accepted bid transfer is $t_1^{\mathrm{LAB}} = 2$, and the auctioneer cannot profitably deviate from the proposed mechanism. In the discriminatory auction, however, awarding $q_1 = 2$ units to bidder 1 always generates more revenue than awarding $q_1 = 1$ unit. Then even when the aggregate bid profile is $(b_1, b_2) = ((2,1),(2,0))$ the auctioneer prefers to allocate to bidder 1 as if bidder 2 bid $b_2' = (0,0)$, and this disclosure policy does not audit the discriminatory auction. In short, the discriminatory auction is not more auditable than the last accepted bid auction because the auctioneer's incentives to mis-allocate are different between the two auction formats.

Essentially the same argument applies to the comparison of the discriminatory and first rejected bid auctions.

**Theorem 4. [PAB audit-incomparable with FRB]** *The discriminatory auction is no more auditable than the first rejected bid uniform price auction and vice versa, $o_{PAB} \not\trianglerighteq o_{FRB}$ and $o_{FRB} \not\trianglerighteq o_{PAB}$.*

In Theorem 3 the case of $Q = 1$ is excluded. When $Q = 1$ the last accepted bid auction is identical to the discriminatory auction, hence Theorem 3 is trivially invalid. However, when $Q = 1$ the first rejected bid auction is a second price auction, and Theorem 4 still applies.

The uniform price auctions also cannot be ordered by auditability. For a uniform price auction to be auditable, the seller must make public the market-clearing price, otherwise idiosyncratic prices could be explicable. In the last accepted bid auction, the seller's incentives are sufficient to ensure that the correct market-clearing price is set (conditional on announcement), since no higher price will clear the market. In the first rejected bid auction, the seller must also attribute the market-clearing price to a particular bidder. If $p^{\star}_{\mathrm{LAB}}(b) > p^{\star}_{\mathrm{FRB}}(b)$ the auctioneer can increase the market-clearing price for all bidders and increase his revenue.[21] However, since information is given in terms of bids it is possible to make public the first rejected bid without revealing the last accepted bid, leaving room for the auctioneer to improve revenue.

**Theorem 5. [LAB audit-incomparable with FRB]** *The last accepted bid and first rejected bid uniform price auctions cannot be ranked in auditability.*

Putting together Theorems 1, 3, 4, and 5 gives the following hierarchy of auditability.

**Proposition 1. [Hierarchy of auditability]** *Let $o_{MENU}$ be any menu mechanism. If $Q > 1$, $o_{MENU} \triangleright o_{PAB}, o_{LAB}, o_{FRB}$, and the latter three cannot be compared. If $Q = 1$, $o_{MENU} \triangleright o_{PAB} = o_{LAB}$, $o_{MENU} \triangleright o_{FRB}$, and $o_{PAB}$ and $o_{LAB}$ cannot be compared to $o_{FRB}$.*

The space of possible multi-unit auction mechanisms is large, and Proposition 1 is an incomplete characterization. We leave a more thorough categorization of auditability to future work.

[21] A similar argument underpins the non-credibility of the second price auction, analyzed in [Akbarpour and Li, 2018].

### 3.2 Outcome-based auditability

Our base model assumes that a disclosure policy partitions the bid space. Because the partition may be freely designed, and the auctioneer faces substantially different incentives under different auction formats, it is difficult to obtain a hierarchy of auditability. In particular, the proofs of the results in Section 3.1 occasionally depend on exotic and unintuitive information structures. To address the possibility that auditability is conceptually biased against mechanism comparison, we now consider disclosure policies which partition the outcome space. Let $\hat{P}(o)$ be given by

$$\hat{P}\left(o\right) = \left\{\mathcal{I}\colon o\left(b\right) = o\left(b'\right) \implies b' \in \mathcal{I}\left(b\right)\right\}.$$

Public information release $\mathcal{I} \in \hat{P}(o)$ may be consistent with multiple outcomes, but all bid profiles consistent with a particular outcome must be contained in the same information set. This restriction on information eliminates the "almost revealing" information structures used to show that $o_{\text{PAB}}$ is not more auditable than $o_{\text{LAB}}$ and $o_{\text{FRB}}$.

As discussed in the introduction and elsewhere, it is impossible to audit the VCG or Spanish auctions with incentive compatible information about outcomes. We prove an intuitively stronger claim, that the first rejected bid auction cannot in general be audited when information must partition the outcome space.[22]

**Theorem 6. [Non-auditability of FRB]** *If $n > 2$ or $Q > 1$, there is no $\mathcal{I} \in \hat{P}(o_{FRB})$ that audits $o_{FRB}$.*

The first rejected bid auction fails to be auditable with disclosure policy $\mathcal{I} \in \hat{P}(o_{\text{FRB}})$ because the auctioneer can misrepresent the market-clearing price without alerting any agents. In a single-unit first rejected bid auction with at least three bidders, the seller can announce the winner's identity and claim any market-clearing price between the highest and second-highest bids. This is consistent with the winning bidder's information, as well as with the losing bidders' information. A similar argument holds when there are two bidders, provided there are at least two units available.

Because outcomes depend on the mechanism implemented, it is not possible to directly compare the auditability of two mechanisms based on outcomes alone. In particular, a public information release that audits one mechanism may not partition the outcome space of another. We therefore define how public information release may implement an outcome partition.

**Definition 6. [Outcome implementation]** Given $\mathcal{O}$, a partition of the outcome space, the disclosure policy $\mathcal{I}$ *implements $\mathcal{O}$ in mechanism $o$* if $o(b)$ and $o(b')$ are in the same partition if and only if $b' \in \mathcal{I}(b)$.

Implementation translates claims on the outcome space to equivalent claims in bid space. Because implementation is defined by "if and only if," given a partition $\mathcal{O}$ and a mechanism $o$ the

[22] The proofs differ, but the fundamental reason for failure — that transfers are based on unobservable opponent bids — is identical across the VCG, Spanish, and first rejected bid uniform price auction formats.

disclosure policy $\mathcal{I}$ that implements $\mathcal{O}$ exists and is unique. Then any claim applying to all disclosure policies $\mathcal{I}$ (subject to implementation) is equivalent to a claim on all outcome partitions $\mathcal{O}$. This allows us to state a version of the hierarchy of auditability for outcome-based information.

**Definition 7.** **[More auditable by outcomes]** If for all partitions of the outcome space $\mathcal{O}$, $\mathcal{I}'$ implements $\mathcal{O}$ in mechanism $o'$ and audits $o'$ whenever $\mathcal{I}$ implements $\mathcal{O}$ in mechanism $o$ and audits $o$, $o'$ is *more auditable by outcomes* than $o$, $o' \trianglerighteq^{\mathrm{OC}} o$.

Implementation is a requirement that information about bids can be translated to information about outcomes realized by the mechanism. In many mechanisms, only a subset of the outcome space may be realized. This gives a degree of freedom in the outcome partition: perhaps one mechanism is not auditable under a given outcome partition because certain seller-preferred outcomes are infeasible in the mechanism. Attention to this form of informational claim establishes partial hierarchy of auditability over common auction formats.

**Proposition 2.** **[Hierarchy of outcome-based auditability]** *Suppose $Q > 1$. Letting $o_{MENU}$ be any menu mechanism,*

$$o_{MENU} \triangleright^{OC} o_{PAB},\ o_{LAB} \triangleright^{OC} o_{FRB}.$$

We have already seen that $o_{\mathrm{MENU}}$ is zero-auditable, and that $o_{\mathrm{FRB}}$ cannot be audited by any $\mathcal{I} \in \hat{P}(o_{\mathrm{FRB}})$. Since the non-informative disclosure $\mathcal{I} \equiv B^n$, corresponding to $\mathcal{O} = \{O\}$, is in $\hat{P}(\cdot)$ and neither $o_{\mathrm{PAB}}$ nor $o_{\mathrm{LAB}}$ is audited by this $\mathcal{I}$, the initial strict ranking in Proposition 2 is immediate. In the appendix, we complete the proof by showing that $o_{\mathrm{PAB}}$ and $o_{\mathrm{LAB}}$ cannot be ranked by auditability in outcomes. The case of $Q = 1$ is excluded since, when a single unit is available, the discriminatory and last accepted bid uniform price auctions are both equivalent to the first price auction.

# 4 Positive marginal costs

The maintained assumption that quantity is freely (and infintely) adjustable is essential to Theorem 2. With nonconstant costs the auctioneer's exogenous incentives may help to audit a mechanism. Suppose that the auctioneer faces cost curve $C : \mathbb{N} \to \mathbb{R}_+$, where $C(Q)$ is the cost of supplying quantity $Q$. For $Q > 0$, let $\Delta C(Q) = C(Q) - C(Q-1)$, and assume that $\Delta C$ is positive and weakly increasing. For the moment, we assume that $C$ is common knowledge. Let $\hat{u}^0$ be the auctioneer's utility net of the cost of supply, $\hat{u}^0(o; C) = u^0(o) - C(\sum_{i=1}^n q_i)$.

**Definition 8.** **[Zero auditable net of costs]** The mechanism $o$ is *zero auditable net of costs $C$* if for any bid profile $b$,

$$o\left(b\right) \in \operatorname*{argmax}_{o' \in \mathcal{X}(b)} \hat{u}^0\left(o'; C\right).$$

A mechanism is zero auditable net of costs if profit maximization on the part of the auctioneer is sufficient for bidders to infer that the mechanism has been run truthfully, even when no public information is released. As mentioned above, Theorem 1 implies that when marginal costs are zero a mechanism is zero auditable net of costs if and only if it is a menu of prices and quantities, or supply is potentially infinite. With nonzero marginal costs, more mechanisms are zero auditable. To show this, we consider mechanisms which are zero auditable but for quantity adjustments, in which aggregate supply is a commitment and the only room for manipulation is in transfers or how fixed aggregate supply is allocated across agents.

**Theorem 7. [Zero auditability with costs]** *Suppose that mechanism $o(\cdot; Q^\star)$ is zero-auditable given supply commitment $Q^\star$, and let $Q(b) = \{Q' \colon \forall i, \exists b_{-i} \text{ s.t. } \sum_{j=1}^{n} q^j(b_i, b_{-i}) = Q'\}$ be the set of feasible common-knowledge quantities given bid profile $b$. Mechanism $o$ is zero-auditable net of costs if and only if for any bid profile $b$, $Q^\star(b) \in \operatorname{argmax}_{Q' \in Q(b)} u^0(o(b; Q')) - C(Q')$.*

*Proof.* Conditional on quantity, the mechanism is zero-auditable. Then if, given bids, quantity selection is optimal the mechanism is zero-auditable net of costs. If quantity selection is not optimal, there is a profit-improving aggregate quantity different from that specified by the mechanism. It follows that the mechanism is not zero-auditable net of costs. □

If a mechanism maximizes the auctioneer's profits conditional on agent actions, it is trivially audited in zero information. The possibility of supply commitment in Theorem 7 is useful in application of the reuslt, and does not represent a fundamental shift in our model. Its role is illustrated in the following corollary.

**Corollary 3. [Zero auditability of common auctions]** *If the auctioneer's marginal cost curve $C$ is common knowledge and aggregate supply is determined so that, for some $p^\star \in [p^\star_{FRB}, p^\star_{LAB}]$, $\Delta C(Q') \leq p^\star$ for $Q' < Q(p^\star)$ and $\Delta C(Q') \geq p^\star$ for $Q' > Q(p^\star)$, the discriminatory auction is zero-auditable. If supply is not determined in this way, the discriminatory auction is not zero-auditable. Regardless of claimed supply, neither uniform price auction format is zero-auditable.*

In a discriminatory auction, a bidder's allocation is sufficient to audit her transfer. Then if the auctioneer may exogenously commit to supply, the discriminatory auction is zero-auditable. With commonly known marginal costs the auctioneer can effectively commit to an elastic supply curve, because ex post profit maximization is consistent with the auctioneer's incentive to misallocate. Corollary 3 is consistent with our initial analysis (in particular, Theorem 2), since under zero marginal costs constraining supply is inconsistent with ex post profit maximization.

The assumption that the auctioneer's marginal costs are common knowledge may be substantially weakened. Even when costs are unknown, the seller can commit to use a mechanism which is profit maximizing. This commitment is zero-auditable, so the discriminatory auction is zero-auditable with essentially zero information about the auctioneer, except that he will select the most profitable allocation. Uniform price auctions are not zero-auditable in this framework, for the same reason that they are not zero-auditable with zero marginal costs: although selecting a profit-maximizing market-clearing price is a believable commitment, without further disclosure the bidders have no mechanism to ensure that they have received identical prices.

Let $\mathcal{C} \subseteq \{C \colon \mathbb{R}_+ \to \mathbb{R}_+\}$ be a set of possible cost curves. A *conditional mechanism* $o : B^n \times \mathcal{C} \to O$ maps submitted bids and a realized cost curve to an outcome. We extend the definition of explicability, so that an outcome is explicable considering costs if is explicable, allowing for the fact that the auctioneer might misrepresent the cost curve $C$.

**Definition 9. [Explicability, considering costs]** Given bid $b_i$ and public information $I$, outcome $o'_i$ is *explicable for bidder* $i$*, considering costs,* if there is $b'_{-i} \in B^{n-1}$ and $C \in \mathcal{C}$ such that $b' = (b_i, b'_{-i}) \in I$ and $o'_i = o^i(b'; C)$. Given bid profile $b$ and public information $I$, outcome $o'$ is *explicable* if $o'_i$ is explicable for each bidder $i$, considering costs, given bid $b_i$ and public information $I$.

As with our initial definition of auditability (and zero-auditability), a mechanism is auditable, considering costs, if its outcomes are weakly revenue maximizing in the set of explicable outcomes, considering costs. A mechanism is zero-auditable, considering costs, if its outcomes are weakly revenue maximizing in the set of explicable outcomes, considering costs, even when no post-allocation disclosure is made.

We now define ex post optimal common auctions, and show that knowledge of profit maximization is sufficient to zero-audit a discriminatory auction, whether or not the cost curve $C$ is known. Ex post optimal common auctions do not differ from their (potentially nonoptimal) counterparts, except that the quantity allocated depends explicitly on both the bids submitted as well as the realized cost curve $C$. In particular,

$$\begin{aligned}
&Q^\star_{\text{PAB}}(b; C) = \max\left\{Q \colon \Delta C(Q) \leq p^\star_{\text{FRB}}(b; Q)\right\};\\
&Q^\star_{\text{LAB}}(b; C) \in \operatorname*{argmax}_{Q} Q p^\star_{\text{LAB}}(b; Q) - C(Q);\\
&Q^\star_{\text{FRB}}(b; C) \in \operatorname*{argmax}_{Q} Q p^\star_{\text{FRB}}(b; Q) - C(Q).
\end{aligned}$$

All ex post optimal mechanisms essentially equate marginal revenue with marginal cost.[23] In the discriminatory auction, this is equivalent to assuming the market-clearing price is (roughly) equal to marginal cost. In the uniform price auctions, however, the market-clearing price may be far from marginal cost; this is familiar from standard monopoly problems.

**Theorem 8. [Zero auditability, considering costs]** *The ex post optimal discriminatory auction is zero-auditable, considering costs. The ex post optimal uniform price auctions are not zero-auditable, considering costs.*

*Proof.* The result for the discriminatory auction follows from Theorem 7. By construction, ex post optimal mechanisms are such that $C \in \operatorname{argmax}_{C' \in \mathcal{C}} u^0\left(o\left(b; C'\right)\right)$. Then it is sufficient to ask whether a mechanism is zero-auditable when the auctioneer announces an optimal quantity commitment between bid submission and allocation. It follows that the discriminatory auction is zero-auditable, considering costs.

[23] The incremental cost used to define $Q^\star_{\text{PAB}}$ addresses the assumption that quantities are discrete.

That the uniform price auctions are not zero-auditable, considering costs, follows from the possibility that the seller assigns heterogeneous market-clearing prices. In the first rejected bid auction it is additionally possible to misrepresent the market-clearing price to all agents. □

Knowledge of profit maximization is sufficient to audit a discriminatory auction, as long as no claims are made regarding the quantity which will arise. Knowledge of profit maximization is not sufficient to audit a uniform price auction. Intuitively, claiming to price discriminate is consistent with revenue maximization, while claiming to offer a consistent market price is inconsistent with revenue maximization, unless additional information is introduced.

### 4.1 Revenue with uncertain costs

Theorem 8 establishes that a seller's claim to maximize revenue is consistent with the discriminatory auction allocation, even when her marginal cost curve is unknown. We now consider the effect of marginal cost uncertainty on revenue in a discriminatory auction. For tractability, we assume now that quantity is real-valued, $Q \in \mathbb{R}_+$,[24] and that bids are decreasing functions, $b : \mathbb{R}_+ \to \mathbb{R}_+$.

There are $n$ bidders, $i \in \{1, \ldots, n\}$, and an auctioneer. All bidders observe signal $s \in [0, 1]$, hidden to the auctioneer. Each bidder's utility is increasing in quantity and quasilinear in transfers, and for each $s$ there is a weakly positive decreasing function $v(\cdot; s) : \mathbb{R}_+ \to \mathbb{R}_+$ so that

$$\bar{u}(q, t; s) = \int_0^q v(x; s)\, dx - t.$$

The bidders' shared type is ordered, so that $q^*(s) \in \operatorname{argmax}_q qv(q; s)$ is single-valued and increasing, and $p^*(s) = v(q^*(s); s)$ is also increasing. The auctioneer has a random marginal cost curve $c \sim F_c$, which is independent of the bidders' signal $s$. The distribution $F_c$ implies a conditional distribution of marginal costs given quantity, $H(p; Q) = \Pr(c(Q) \leq p)$. Under the assumption that each marginal cost curve $c \in \operatorname{Supp} F_c$ is weakly increasing, $H(p; Q) \succeq_{\text{FOSD}} H(p; Q')$ whenever $Q \geq Q'$.

Bidders participate in a discriminatory auction, where the quantity sold maximizes profits conditional on realized bids and marginal costs; if inverse bids are $\varphi^i$, the realized quantity is

$$Q^\star = \sum_{i=1}^n \varphi^i(p^\star), \quad p^\star \in \operatorname*{argmax}_p \sum_{i=1}^n \int_0^{\varphi^i(p)} b^i(x)\, dx - \int_0^{\sum_{i=1}^n \varphi^i(p)} c(x)\, dx.$$

We consider Bayesian Nash equilibria $(b^i)_{i=1}^n$ of this game, in which each bidder $i$'s bidding strategy $b^i$ is a best response to her opponents' bidding strategies $b^{-i}$. If $c$ is random, auction revenue is strictly lower than if an optimal supply curve is announced prior to the auction.

**Proposition 3. [Non-auditability reduces revenue]** *In a discriminatory auction, equilibrium revenue is (strictly) lower when bidders believe quantity is determined by profit maximization conditional on (random) marginal costs than when an auditable deterministic supply curve is announced prior to the auction.*

[24] All results hold as the distribution over feasible quantities places all weight on integer-valued allocations.

*Proof.* Let $G^i(q;b)$ be the probability that bidder $i$ recieves quantity $q_i \leq q$, conditional on submitting bid function $b$. Note that $G^i$ is also determined by the auctioneer's profit maximization problem, and by opponent bidding strategies. Standard integration by parts gives that bidder $i$'s expected utility is

$$u\left(b^i, b^{-i}; s\right) = \int_0^\infty \left(v\left(x;s\right) - b^i\left(x\right)\right)\left(1 - G^i\left(x;b^i\right)\right) dx.$$

Integration by parts and the calculus of variations together give the agent's first order condition for optimality,

$$-\left(v\left(q;s\right) - b^i\left(q\right)\right) G_b^i\left(q;b^i\right) = 1 - G^i\left(q;b^i\right). \tag{1}$$

In the discriminatory auction, it is optimal for the auctioneer to award all units with a bid below marginal cost. [Pycia and Woodward, 2019] show that the bidding equilibrium must be symmetric; then supposing that all opponents $-i$ use bid function $b$,

$$\begin{aligned}1 - G^i\left(q;b^i\right) = \Pr\left(q_i \geq q | b^i\right) &= \Pr\left(c\left(q + (n-1)\varphi\left(b^i\left(q\right)\right)\right) \leq b^i\left(q\right)\right) \\ &= H\left(b^i\left(q\right); q + (n-1)\varphi\left(b^i\left(q\right)\right)\right).\end{aligned}$$

Then we have

$$\begin{aligned}G_b^i\left(q;b^i\right) = &-H_p\left(b^i\left(q\right); q + (n-1)\varphi\left(b^i\left(q\right)\right)\right) \\ &- (n-1)\varphi_p\left(b^i\left(q\right)\right) H_Q\left(b^i\left(q\right); q + (n-1)\varphi\left(b^i\left(q\right)\right)\right).\end{aligned}$$

Since equilibrium is symmetric, we have $\varphi(b^i(q)) = q$ in equilibrium. Then in equilibrium condition (1) becomes

$$\left(v\left(q;s\right) - b\left(q\right)\right)\left(H_p\left(b\left(q\right);nq\right) + (n-1)\varphi_p\left(b\left(q\right)\right) H_Q\left(b\left(q\right);nq\right)\right) = H\left(b\left(q\right);nq\right).$$

This is exactly the first order condition given in [Pycia and Woodward, 2019] for equilibrium bids given random supply and reserve. Since [Pycia and Woodward, 2019] show that (potentially correlated) random supply and reserve generate less revenue than deterministic supply and reserve, the result follows. □

When supply is deterministic, bids in the discriminatory auction are perfectly flat until the per capita maximum quantity. Determining an optimal supply curve is equivalent to solving the monopolist's problem for any realization of bidder type $s$. The marginal costs of increased allocation are not just the seller's marginal cost of increasing $Q$, but must also take into account that increasing $Q$ lowers submitted bids. Then the marginal revenue associated with allocating additional quantity (and price discriminating) is above the seller's marginal cost of this increase, and even given an optimal supply curve, the auctioneer can still improve revenues by increasing allocation ex post. It follows that discriminatory auctions with optimal supply curves are not zero-auditable, even though ex post optimal discriminatory auctions are zero-auditable. By corollary, bids will be relatively low if there is no disclosure. We conclude that elastic supply commitment, via a disclosure policy that audits an auction's outcomes, has positive value to the auctioneer. When marginal cost curves $c$ are random, the value of commitment is strictly positive.

With regard to the uniform price auction, note first that true marginal values are an upper bound for bids in a uniform price auction, regardless of the underlying distribution of quantity and its relationship to bids. Then since the deterministic discriminatory auction generates higher revenue than any stochastic-outcome uniform price auction with truthful bids [Pycia and Woodward, 2019], and the two auction formats are revenue equivalent when supply is optimized, it follows that auditability also has value in a uniform price auction. This comparison holds even though the uniform price auction is not zero-auditable under quantity optimization.[25]

# 5 Conclusion

In this paper we have introduced the concept of auditability. A disclosure policy audits an auction's outcome if each bidder's outcome is consistent with the assumption that the seller is maximizing profits, conditional on consistency with public and bidder-private information. We show that a menu can be audited with no information while common multi-unit auction formats require at least some information to audit their outcomes. Without further restrictions the relative auditability of common multi-unit auction formats is ambiguous.

We show that when the seller is permitted to condition aggregate supply on realized bids, the discriminatory auction becomes auditable in zero information. This is not the case for uniform price auctions. However, zero-auditability under adjustable supply holds only if the supply curve is the seller's true supply curve (his inverse marginal cost). Because equilibrium bids depend on claimed supply, commitment to a supply curve which differs from the seller's true supply curve can improve the auctioneer's expected revenue. It follows that even when a mechanism can be zero-audited, nontrivial information disclosure may improve the mechanism's revenue.

# References


Atila Abdulkadiroğlu and Tayfun Sönmez. School choice: A mechanism design approach. *American economic review*, 93(3):729–747, 2003.

Mohammad Akbarpour and Shengwu Li. Credible mechanisms. 2018.

Jose Alcalde and Antonio Romero-Medina. Strategy-proof fair school placement. 2015.


[25] [Woodward, 2019] shows that all pure strategy equilibria are symmetric in this model of the uniform price auction. This implies an equilibrium ability to audit auction outcomes: because equilibrium strategies are symmetric, the auctioneer's incentive to alter an agent's allocation is identical for all agents, and each agent can infer her opponents' outcomes from her own. This is possible due to the lack of private information in this auction model. Because the revenue comparison in [Pycia and Woodward, 2019] is strict, Proposition 3 is valid in the presence of at least some private information.

Olivier Armantier and Erwann Sbaï. Estimation and comparison of treasury auction formats when bidders are asymmetric. *Journal of Applied Econometrics*, 21(6):745–779, 2006. ISSN 1099-1255. doi: 10.1002/jae.875.

Olivier Armantier and Erwann Sbaï. Comparison of alternative payment mechanisms for french treasury auctions. *Annals of Economics and Statistics/Annales d'Économie et de Statistique*, pages 135–160, 2009.

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.

Elchanan Ben-Porath and Barton L Lipman. Implementation with partial provability. *Journal of Economic Theory*, 147(5):1689–1724, 2012.

Helmut Bester and Roland Strausz. Imperfect commitment and the revelation principle: the multi-agent case. *Economics Letters*, 69(2):165–171, 2000.

Helmut Bester and Roland Strausz. Contracting with imperfect commitment and the revelation principle: the single agent case. *Econometrica*, 69(4):1077–1098, 2001.

Sayan Bhattacharya, Elias Koutsoupias, Janardhan Kulkarni, Stefano Leonardi, Tim Roughgarden, and Xiaoming Xu. Near-optimal multi-unit auctions with ordered bidders. In *Proceedings of the fourteenth ACM conference on Electronic commerce*, pages 91–102. ACM, 2013.

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

Jesse Bull and Joel Watson. Hard evidence and mechanism design. *Games and Economic Behavior*, 58(1):75–93, 2007.

Justin Burkett and Kyle Woodward. Reserve prices eliminate low revenue equilibria in uniform price auctions. 2020a.

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

Giacomo Calzolari and Alessandro Pavan. On the optimality of privacy in sequential contracting. *Journal of Economic theory*, 130(1):168–204, 2006.

David Cantala and Szilvia Pápai. Reasonably and securely stable matching. Technical report, 2014.

Sara Castellanos and Marco Oviedo. Optimal bidding in the mexican treasury securities primary auctions: results of a structural econometric approach. *Cuadernos de economía*, 45(131):3–28, 2008.

Raymond Deneckere and Sergei Severinov. Mechanism design with partial state verifiability. *Games and Economic Behavior*, 64(2):487–513, 2008.

Lars Ehlers and Thayer Morrill. (il)legal assignments in school choice. 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.

Andrew V Goldberg, Jason D Hartline, Anna R Karlin, Michael Saks, and Andrew Wright. Competitive auctions. *Games and Economic Behavior*, 55(2):242–269, 2006.

Pär Holmberg. Supply function equilibria of pay-as-bid auctions. *Journal of Regulatory Economics*, 36(2):154–177, 2009.

Ali Hortaçsu and Jakub Kastl. Valuing dealers' informational advantage: a study of canadian treasury auctions. *Econometrica*, 80(6):2511–2542, 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, Jakub Kastl, and Allen Zhang. Bid shading and bidder surplus in the us treasury auction system. *American Economic Review*, 108(1):147–69, 2018.

Boo-Sung Kang and Steven L Puller. The effect of auction format on efficiency and revenue in divisible good auctions: A test using Korean Treasury auctions. *The Journal of Industrial Economics*, 56(2):290–332, 2008.

Marco LiCalzi and Alessandro Pavan. Tilting the supply schedule to enhance competition in uniform-price auctions. *European Economic Review*, 49(1):227–250, 2005.

David McAdams. Adjustable supply in uniform price auctions: Non-commitment as a strategic tool. *Economics Letters*, 95(1):48–53, 2007.

David McAdams and Michael Schwarz. Credible sales mechanisms and intermediaries. *American Economic Review*, 97(1):260–276, 2007a.

David McAdams and Michael Schwarz. Who pays when auction rules are bent? *International Journal of Industrial Organization*, 25(5):1144–1157, 2007b.

R Preston McAfee and Daniel Vincent. Sequentially optimal auctions. *Games and Economic Behavior*, 18(2):246–276, 1997.

Dilip Mookherjee and Ivan Png. Optimal auditing, insurance, and redistribution. *The Quarterly Journal of Economics*, 104(2):399–415, 1989.

Tymofiy Mylovanov and Andriy Zapechelnyuk. Optimal allocation with ex post verification and limited penalties. *American Economic Review*, 107(9):2666–94, 2017.

Andrew Postlewaite and David Schmeidler. Implementation in differential information economies. *Journal of Economic Theory*, 39(1):14–33, 1986.

Marek Pycia and Kyle Woodward. Auctions of homogeneous goods: A case for pay-as-bid. 2019.

Vasiliki Skreta. Sequentially optimal mechanisms. *The Review of Economic Studies*, 73(4):1085–1111, 2006.

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

Vasiliki Skreta. Optimal auction design under non-commitment. *Journal of Economic Theory*, 159: 854–890, 2015.

Robert M Townsend. Optimal contracts and competitive markets with costly state verification. *Journal of Economic theory*, 21(2):265–293, 1979.

Peter Troyan, David Delacretaz, and Andrew Kloosterman. Essentially stable matchings. 2018.

Hannu Vartiainen. Auction design without commitment. *Journal of the European Economic Association*, 11(2):316–342, 2013.

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

Kyle Woodward. Hybrid mechanisms for the sale of divisible goods. 2019.

# A Proofs

*Proof of Theorem 3.* We provide a disclosure policy that audits $o_{\text{PAB}}$ but not $o_{\text{LAB}}$, then a disclosure policy that audits $o_{\text{LAB}}$ but not $o_{\text{PAB}}$.

We first show that $o_{\text{LAB}} \not\trianglerighteq o_{\text{PAB}}$. Let $\mathcal{I}$ fully reveal the aggregate allocation, $\mathcal{I}(b) = \{b' \colon q(b') = q(b)\}$. Then for all $b' \in \mathcal{I}(b)$, $u^0(o_{\text{PAB}}(b')) = u^0(o_{\text{PAB}}(b))$, and selecting a different allocation weakly reduces the auctioneer's revenue. Then $\mathcal{I}$ audits $o_{\text{PAB}}$. Now, fix bidders $i$ and $j \neq i$. For any $p_i$ and $p_j \leq p_i$ define bids by

$$b_{iq}(p_i, p_j) = \begin{cases} p_i & \text{if } q \leq Q-1, \\ 0 & \text{otherwise;} \end{cases} \qquad b_{jq}(p_i, p_j) = \begin{cases} p_j & \text{if } q = 1, \\ 0 & \text{otherwise.} \end{cases}$$

Suppose that $b_{kq} = 0$ for all $q$ and all $k \neq i, j$. Then $q^i(b(p_i, p_j)) = Q - 1$, $q^j(b(p_i, p_j)) = 1$, and $q^k(b(p_i, p_j)) = 0$ for all $k \neq i, j$. Note that $p^\star_{\text{LAB}}(b(p_i, p_j)) = p_j$; since the allocation is fixed, $b(p_i, p'_j) \in \mathcal{I}(b(p_i, p_j))$ for all $p'_j \in (0, p_i]$. Letting $p'_j > p_j$ without loss of generality, $t^i_{\text{LAB}}(b(p_i, p_j)) = (Q-1)p_j < (Q-1)p'_j = t^i_{\text{LAB}}(b(p_i, p'_j))$, and $\mathcal{I}$ does not audit $o_{\text{LAB}}$.

We now show $o_{\text{PAB}} \not\trianglerighteq o_{\text{LAB}}$. For two agents $i$ and $j$, $i \neq j$, let $\tilde{q}_i = 1$ and $\tilde{q}_j = Q - 1$, and let $\tilde{q}_\ell = 0$ for all $\ell \neq i, j$. Letting $p > 0$, consider two bid profiles $b$ and $b'$,

$$b_{\ell k} = \begin{cases} p & \text{if } k \leq \tilde{q}_\ell, \\ \left(\frac{Q-1}{Q}\right) p & \text{if } k = Q \text{ and } \ell = j, \\ 0 & \text{otherwise;} \end{cases} \quad b'_{\ell k} = \begin{cases} 0 & \text{if } \ell = i \text{ and } k = \tilde{q}_\ell, \\ b_{\ell k} & \text{otherwise.} \end{cases}$$

Let the disclosure policy $\mathcal{I}$ fully reveal the aggregate bid profile, as long as it is not $b$ or $b'$,

$$\mathcal{I}\left(\tilde{b}\right) = \begin{cases} \left\{\tilde{b}\right\} & \text{if } \tilde{b} \notin \{b, b'\}, \\ \{b, b'\} & \text{otherwise.} \end{cases}$$

Because $\mathcal{I}$ is single-valued for all $\tilde{b} \notin \{b, b'\}$, the auctioneer can potentially misrepresent outcomes only when $\tilde{b} \in \{b, b'\}$. Note that $q^i(b) = 1$ and $q^j(b) = Q - 1$, while $q^i(b') = 0$ and $q^j(b') = Q$; $q^\ell(b) = q^\ell(b') = 0$ for all $\ell \neq i, j$. Furthermore, $p^\star_{\text{LAB}}(b) = p$ and $p^\star_{\text{LAB}}(b') = (Q-1)p/Q$. Then $t^j_{\text{LAB}}(b) = t^j_{\text{LAB}}(b')$. Conditional on $\tilde{b} \in \{b, b'\}$, the outcome is completely determined by agent $i$'s bid, so the auctioneer cannot mis-allocate to agent $i$. It follows that $\mathcal{I}$ audits $o_{\text{LAB}}$. However, given bid profile $b$ the outcome $o_j = (Q, (Q^2-1)p/Q)$ is explicable for agent $j$ in $o_{\text{PAB}}$, and generates strictly more revenue than the correct outcome $o^j(b) = (Q-1, (Q-1)p)$. Then $\mathcal{I}$ does not audit $o_{\text{PAB}}$. $\square$

*Proof of Theorem 4.* Showing $o_{\text{FRB}} \not\trianglerighteq o_{\text{PAB}}$ is essentially identical to the proof that $o_{\text{LAB}} \not\trianglerighteq o_{\text{PAB}}$ in Theorem 3, and is omitted here.

We now show $o_{\text{PAB}} \not\trianglerighteq o_{\text{FRB}}$. This proof is similar to the proof that $o_{\text{PAB}} \not\trianglerighteq o_{\text{LAB}}$, but the bids used to construct a disclosure policy are distinct. For two agents $i$ and $j$, $i \neq j$, let $\tilde{q}_i = 1$ and

$\tilde{q}_j = Q-1$, and let $\tilde{q}_\ell = 0$ for all $\ell \neq i,j$. Letting $5p > 0$, consider two bid profiles $b$ and $b'$,

$$b_{\ell k} = \begin{cases} 2p & \text{if } k \leq \tilde{q}_\ell, \\ p & \text{if } k = \tilde{q}_\ell \text{ and } \ell \in \{i,j\}, \\ 0 & \text{otherwise;} \end{cases} \quad b'_{\ell k} = \begin{cases} 2p & \text{if } k \leq \tilde{q}_\ell \text{ and } \ell = j, \\ p & \text{if } k = \tilde{q}_\ell + 1 \text{ and } \ell = j, \\ \left(\frac{Q-1}{Q}\right) p & \text{if } k \leq \tilde{q}_\ell + 1 \text{ and } \ell = i, \\ 0 & \text{otherwise.} \end{cases}$$

Let the disclosure policy $\mathcal{I}$ fully reveal the aggregate bid profile, as long as it is not $b$ or $b'$,

$$\mathcal{I}\left(\tilde{b}\right) = \begin{cases} \left\{\tilde{b}\right\} & \text{if } \tilde{b} \notin \{b, b'\}, \\ \{b, b'\} & \text{otherwise.} \end{cases}$$

Because $\mathcal{I}$ is single-valued for all $\tilde{b} \notin \{b,b'\}$, the auctioneer can potentially misrepresent outcomes only when $\tilde{b} \in \{b,b'\}$. Note that $q^i(b) = 1$ and $q^j(b) = Q-1$, whilc $q^i(b') = 0$ and $q^j(b') = Q$; $q^\ell(b) = q^\ell(b') = 0$ for all $\ell \neq i,j$. Furthermore, $p^\star_{\text{FRB}}(b) = p$ and $p^\star_{\text{FRB}}(b') = (Q-1)p/Q$. Then $t^j_{\text{FRB}}(b) = t^j_{\text{FRB}}(b')$. Conditional on $\tilde{b} \in \{b,b'\}$, the outcome is completely determined by agent $i$'s bid, so the auctioneer cannot mis-allocate to agent $i$. It follows that $\mathcal{I}$ audits $o_{\text{FRB}}$. However, given bid profile $b$ the outcome $o_j = (Q, (2Q-1)p)$ is explicable for agent $j$ in $o_{\text{PAB}}$, and generates strictly more revenue than the correct outcome $o^j(b) = (Q-1, 2(Q-1)p)$. Then $\mathcal{I}$ does not audit $o_{\text{PAB}}$. $\square$

*Proof of Theorem 5.* We first show $o_{\text{FRB}} \not\trianglerighteq o_{\text{LAB}}$. Consider the disclosure policy $\mathcal{I}(b) = \{b' \colon q(b') = q(b) \text{ and } p^\star_{\text{LAB}}(b') = p^\star_{\text{LAB}}(b)\}$. Because the auctioneer cannot misrepresent the allocation $q$ or the market-clearing price $p^\star_{\text{LAB}}$, $\mathcal{I}$ audits $o_{\text{LAB}}$.[26] Fix an allocation $\tilde{q}$ such that $\sum_{i=1}^n \tilde{q}_i = Q$, and let agent $i$ be such that $\tilde{q}_i < Q$. Letting $2p > 0$, consider two bid profiles $b$ and $b'$,

$$b_{jk} = \begin{cases} 2p & \text{if } k \leq \tilde{q}_j, \\ 0 & \text{otherwise;} \end{cases} \quad b'_{jk} = \begin{cases} 2p & \text{if } k \leq \tilde{q}_j, \\ p & \text{if } k = \tilde{q}_j + 1 \text{ and } j = i, \\ 0 & \text{otherwise.} \end{cases}$$

Note that $p^\star_{\text{LAB}}(b) = p^\star_{\text{LAB}}(b')$, and $q(b) = q(b') = \tilde{q}$. Then $b' \in \mathcal{I}(b)$. Because $\tilde{q}_i < Q$, there is an agent $j \neq i$ such that $\tilde{q}_j \geq 1$. Although $p^\star_{\text{LAB}}(b) = p^\star_{\text{LAB}}(b')$, $p^\star_{\text{FRB}}(b) = 0 < p = p^\star_{\text{FRB}}(b')$. Then the individual outcome $o_j = (\tilde{q}_j, p\tilde{q}_j)$ is explicable for bidder $j$ under $o_{\text{FRB}}$, even though the correct outcome is $o_j = (\tilde{q}_j, 0)$. Then $\mathcal{I}$ does not audit $o_{\text{FRB}}$, and $o_{\text{FRB}} \not\trianglerighteq o_{\text{LAB}}$.

Showing that $o_{\text{LAB}} \not\trianglerighteq o_{\text{FRB}}$ is essentially identical to the proof that $o_{\text{FRB}} \not\trianglerighteq o_{\text{LAB}}$, considering instead $\mathcal{I}(b) = \{b' \colon q(b') = q(b),\ p^\star_{\text{FRB}}(b') = p^\star_{\text{FRB}}(b), \text{ and } S(b') = S(b)\}$, where $S(b) = \min\{i \colon b^i_{q^i(b)+1} = p^\star_{\text{FRB}}(b)\}$ identifies a price-setting bidder. $\square$

[26] Announcing $I \neq \mathcal{I}(b)$ cannot increase the aggregate quantity supplied, but can affect the market-clearing price. However, the market-clearing price can never exceed $p^\star_{\text{LAB}}$, so there is no profitable misallocation.

*Proof of Theorem 6.* We show first that $o_{\text{FRB}}$ cannot be audited (in $\hat{P}$) when a single unit is for sale, provided there are at least $n = 3$ bidders. For simplicity of exposition, re-parameterize an outcome $o = (i, t)$ as the identity of the winning bidder and their transfer to the seller, and let the bid profile $b$ be an $n$-dimensional vector of unit bids.[27] Let $b = (2p, 0, \ldots, 0)$. Then $o_{\text{FRB}}(b) = (1, 0)$, since bidder 1 submits the highest bid and all other bids are 0. Because all disclosure policies $\mathcal{I} \in \hat{P}$ give information only about outcomes, the outcome $o' = (1, p)$ is explicable for all agents: bidder 1 believes that some agent $i \neq 1$ has bid $p$, and all bidders $i \neq 1$ believe there is some agent $j_i \neq 1, i$ that has bid $p$. The auctioneer's transfer is greater in outcome $o'$ than in outcome $o$, hence $o_{\text{FRB}}$ cannot be audited from $\hat{P}$ when $Q = 1$ and $n \geq 3$.

Now suppose that there are at least $n = 2$ bidders, and $Q \geq 2$. Let $\tilde{q}_i, \tilde{q}_j > 0$ be such that $\tilde{q}_i + \tilde{q}_j = Q$. Let $2p > 0$ and consider three bid profiles, $b$, $b'^i$, and $b'^j$,

$$b_{\ell k} = \begin{cases} 2p & \text{if } \ell \in \{i, j\} \text{ and } k \leq \tilde{q}_\ell, \\ 0 & \text{otherwise;} \end{cases}$$

$$b'^i_{\ell k} = \begin{cases} 2p & \text{if } \ell \in \{i, j\} \text{ and } k \leq \tilde{q}_\ell, \\ p & \text{if } \ell = i \text{ and } k = \tilde{q}_\ell + 1, \\ 0 & \text{otherwise;} \end{cases} \qquad b'^j_{\ell k} = \begin{cases} 2p & \text{if } \ell \in \{i, j\} \text{ and} k \leq \tilde{q}_\ell, \\ p & \text{if } \ell = j \text{ and } k = \tilde{q}_\ell + 1, \\ 0 & \text{otherwise.} \end{cases}$$

Note that $q(b) = q(b'^i) = q(b'^j)$, and the allocation is $\tilde{q} = (\tilde{q}_i, \tilde{q}_j, 0, \ldots, 0)$ in case. However, $p^\star_{\text{FRB}}(b) = 0$ and $p^\star_{\text{FRB}}(b'^i) = p^\star_{\text{FRB}}(b'^j) = p$. Because $o(b'^i) = o(b'^j)$ is explicable given bid profile $b$ and any information $\mathcal{I}(b'^i)$, $o_{\text{FRB}}$ cannot be audited by any $\mathcal{I} \in \hat{P}$. □

*Proof of Proposition 2.* It remains to be seen that $o_{\text{PAB}}$ and $o_{\text{LAB}}$ can be audited with outcome-based information, and cannot be ranked within this set.

First, a disclosure policy that announces the allocation profile is outcome-based and audits $o_{\text{PAB}}$ (Lemma 1). We then need only to provide a disclosure policy $\mathcal{I} \in \hat{P}$ that audits $o_{\text{LAB}}$ but not $o_{\text{PAB}}$. Fix $\hat{p} > 0$ and form a partition of the outcome space $\mathcal{O} = \hat{\mathcal{O}} \cup \mathcal{O}^x \cup \{(\cup_{O \in \hat{\mathcal{O}} \cup \mathcal{O}^x} O)^C\}$,

$$\hat{\mathcal{O}} = \left\{ \{((q_\ell, p q_\ell))_{\ell=1}^n\} : p \in \mathbb{R}_+, \sum_{\ell=1}^n q_\ell = Q \right\},$$
$$\mathcal{O}^x = \{\{((Q-1, 2(Q-1)\hat{p}), (1, 4\hat{p})), ((Q, (2Q-1)\hat{p}), (0,0))\} \times (\times_{\ell > 2} \{(0,0)\})\}.$$

The partition $\mathcal{O}$ corresponds to bidders 1 and 2 splitting the market at any price, or one of the two outcomes in $\mathcal{O}^x$ arising. Note that $o_1^x = (Q-1, 2(Q-1)\hat{p})$ and $o_2^x = (1, 4\hat{p})$ is not a feasible outcome in $o_{\text{LAB}}$, as the two bidders are paying different per-unit prices. Let $B^{12}(b) = \{b' : p^\star_{\text{LAB}}(b') = p^\star_{\text{LAB}}(b) \text{ and } \forall \ell \notin \{1, 2\}, q^\ell(b') = 0\}$. Then if $\mathcal{I}_{\text{LAB}}$ implements $\mathcal{O}$ in $o_{\text{LAB}}$,

$$\mathcal{I}(b) = \{b' : p^\star_{\text{LAB}}(b') = p^\star_{\text{LAB}}(b) \text{ and } q(b') = q(b)\}.$$

[27] In the case of a single-unit auction, this is sufficient.

By previous arguments, $\mathcal{I}_{\text{LAB}}$ audits $o_{\text{LAB}}$. Although allocations are identical across auction mechanisms, it will generally not be the case that discriminatory auction transfers equal last accepted bid auction transfers, and the $\mathcal{I}_{\text{PAB}}$ that implements $\mathcal{O}$ in $o_{\text{PAB}}$ will differ from $\mathcal{I}_{\text{LAB}}$. Consider two bid profiles, $b$ and $b'$,

$$b_{\ell k} = \begin{cases} 2\hat{p} & \text{if } \ell = 1, \\ 4\hat{p} & \text{if } \ell = 2 \text{ and } k = 1, \\ 0 & \text{otherwise;} \end{cases} \quad b'_{\ell k} = \begin{cases} 2\hat{p} & \text{if } \ell = 1, \\ 0 & \text{otherwise.} \end{cases}$$

Then $o^1_{\text{PAB}}(b) = (Q-1, 2(Q-1)\hat{p})$ and $o^2_{\text{PAB}} = (1, 4\hat{p})$, and $o^1_{\text{PAB}}(b') = (Q, 2Q\hat{p})$ and $o^2_{\text{PAB}}(b') = (0,0)$. Then if $\mathcal{I}_{\text{PAB}}$ implements $\mathcal{O}$ in $o_{\text{PAB}}$, $b' \in \mathcal{I}_{\text{PAB}}(b)$ . Following the arguments in Theorem 3 it follows that $\mathcal{I}_{\text{LAB}}$ does not audit $o_{\text{PAB}}$. $\square$