Self-Auditable Auctions

Kyle Woodward*

January 2019

Abstract

We consider the amount of information necessary to verify that an auction has been run according to the specified rules. A mechanism is *audited* by a post-allocation information release if each outcome maximizes the auctioneer's utility, conditional on consistency with the information released. One mechanism is more auditable than another if any information that audits the latter also audits the former. In many multi-unit auctions the number of units supplied is not a priori constrained to equal the claimed supply. When the seller cannot commit to any bounds on supply, only menus are auditable without additional information. Generally, menus are more auditable than discriminatory auctions, which are more auditable than uniform price auctions. Provided the type space is sufficiently rich, this ordering holds even if participants are only suspicious of unexpectedly utility-negative outcomes. Constraining the space of information release strategies to those observed in practice yields the same hierarchy. In line with previous work, we find empirical support for this hierarchy of auditability, and show that perceived corruption is positively correlated with auditable sales of sovereign debt.

1 Introduction

In multi-unit auctions for homogeneous goods, buyers have reason to be skeptical of not only their payments but also their allocations. After bids are submitted a seller may arbitrarily select a quantity to supply, potentially improving his profits over the claimed mechanism. Even if a bidder's payment is verifiable conditional on her allocated quantity, it may be that the quantity was not honestly determined.

*UNC Chapel Hill; kyle.woodward@unc.edu. We would like to thank Gary Biglaiser, Justin Burkett, Zhiguo He, Fei Li, Shengwu Li, Marek Pycia, and Larry Samuelson for valuable discussions and feedback.

Arguably the chief application of multi-unit auctions is the sale of sovereign debt. In 2018, for example, the United States sold over $10 trillion in Treasury securities via a uniform price auction. Crucially, in many cases there is no natural constraint on the number of treasury securities available. In the presence of favorable bids, a treasury could accept more bids than advertised, without informing bidders of this action. If this is a realistic possibility, auction participants may adjust their bids accordingly. In this paper, we address the question of how much information is necessary to certify that the auction was run according to the stated rules, and thus ensure bidders that they are playing the game they are promised.

We consider the possibility of post-auction information release, under the constraint that information must be consistent with each individual bidder's information set. For example, the transfer to the auctioneer in a discriminatory auction is verifiable, since bidders can plainly see whether they are paying what they bid for the quantities they receive, so to (fully) verify the discriminatory auction it is only necessary to show that the proper quantity was sold. The auctioneer can, after the auction is run, make public the identities of the bidders and the quantities they received. Each bidder can verify that they received the quantity the auctioneer claims they received, and that the total quantity allocated is equal to the total quantity offered. We say that information sufficient to determine the aggregate quantity awarded *audits* the discriminatory auction.

A practical example solidifies this concept. U.S. Treasury securities are sold through uniform-price auctions where the per-unit price is equal to the highest market clearing price (the last accepted bid). Following the auction the Treasury makes public the market clearing price. Each bidder can easily verify that the amount paid is equal to the quantity received, times the market clearing price. Without public announcement of the market clearing price personal prices could be assigned, where individual market clearing prices are the last bid accepted from the particular bidder. This would weakly improve seller revenue in all circumstances, and strictly improve seller revenue in most circumstances. Public announcement of the market clearing price can be viewed as a commitment device to ensure that the auction was run honestly.

In our model the structure of information release is on bidders' information about actions taken, and not on the manner in which information is made public nor on how participants derive inferences about outcomes. Although it is intuitively simpler to make claims about what kind of information is released (see the statement above about quantity information auditing the discriminatory auction),

it is mathematically more straightforward to work directly with what agents know about bid profiles. Practically, it may be natural to constrain our notion of information release to lie within a particular set, for example the set of information structures that determine the market-clearing price in an auction.[1] 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.[2]

Our central result uses public information release to construct an auditability hierarchy of multi-unit auctions. One mechanism is more auditable than another if any public information release that verifies the latter also verfies the former. In this construction posted price mechanisms are zero-auditable (auditable without any additional information), and are more auditable than discriminatory auctions, which are more auditable than uniform-price auctions. Posted price mechanisms are zero-auditable because any participant can see whether the quantity and price they receive corresponds to what they offered to purchase. Since they are auditable, they are more auditable than any other mechanism. Intuitively, for a discriminatory auction to be auditable it is necessary to publicly announce that the quantity sold is equal to the quantity offered.[3] For a uniform price auction to be auditable, it is necessary to publicly announce that the quantity sold is equal to the quantity offered, and that each bidder is paying the same market clearing price per unit.[4] Since, to be auditable, the uniform price auction must verify that the correct quantity was sold, which is the only information necessary to verify the discriminatory auction, the discriminatory auction is more auditable than the uniform-price auction.

Auditability is a companion to credibility, as defined in Akbarpour and Li [2018]; as a concept

[1] Public information regarding U.S. Treasury competitive bids falls into this class.

[2] Under this interpretation it is natural to ask whether reporting misinformation is incentive compatible for the bidder. We show that if 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, collusion between a bidder and the auctioneer is possible only if the auctioneer can conditionally deviate from the stated mechanism, using knowledge of bidder preferences to determine when it is appropriate to do so.

[3] 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, since we take as given that the auctioneer claims to sell via a particular mechanism (as observed in practice) and address the believability of this claim, while we do not consider the mechanism design question.

[4] Depending on the market clearing price, it may also be necessary to certify that the correct price was selected. Thus there is a hierarchy of credibility even within the class of uniform price auctions.

it is neither weaker nor stronger. Auditability and credibility share the common thread of inference from what the auctioneer would like to achieve. In both cases, market participants consider whether the seller could improve his outcome by breaking the rules of the mechanism while remaining consistent with the participants' own information. To the extent auditability is stronger, it is because auditability is defined as a property on all action profiles while credibility is defined on equilibrium strategies. The benefit of a stronger requirement is a stronger set of results. The main result of this paper is a hierarchy of auditability in auctions; this is not possible with credibility, which is a binary concept. In the context of multi-unit auctions a property which holds for all action profiles is useful for comparing information derived from submitted bids. It is known (Ausubel et al. [2014], Burkett and Woodward [2018], and others) that equilibrium bids differ significantly between multi-unit auction formats; in fact, the set of equilibrium bids may have very little overlap.[5] Then verifying one auction's outcome with bid information may have no implications for another auction's outcome. Nonetheless the "for all actions" requirement has the natural interpretation that information policies may persist across time, as model fundamentals change. As long as the set of available bids is unchanging, information will continue auditing a particular auction format.

We derive our results for static deterministic mechanisms with a fixed bid space, and assume the auctioneer is interested only in revenue. Restricting attention to static mechanisms covers a substantial portion of multi-unit auctions implemented in practice, and greatly simplifies our results. That the auctioneer is concerned only (or primarily) with his own revenue is a familiar assumption from the auctions literature. If the auctioneer is additionally concerned with, for example, who gets what, it becomes more difficult to audit an auction outcome. Relaxing these two assumptions is left to future work.

Deterministic supply and mechanism-independent feasible bids are central to our results, and are also central to the concept of auditability. With nondeterministic mechanisms it may be impossible to audit any one-shot outcome: the auctioneer may simply claim a favorable random draw.[6] It would be natural to define auditability as testing the claimed distribution of randomness against

[5]This is in contrast with single-unit auctions, in which bids are real numbers. In reasonable models the set of equilibrium bids in a first price auction is a subset of the set of bids in the truth-telling equilibrium of the second price auction.

[6]This is not true of all nondeterministic mechanisms. For example, in a discriminatory auction where only the tiebreaking rule is nondeterministic, quantity remains sufficient to audit outcomes. If the aggregate quantity available is also random, the discriminatory auction cannot be audited by any information about bids.

the observed distribution of randomness over time, but we do not do so here. Fixing the bid space across auctions is essential to the definition of auditability: releasing information about bids depends on the bid space, and if bid spaces change between mechanisms, "the same information release" may not make sense across two formats (this is related to the above comments on equilibrium inference between auction formats).

Our theoretical results give an intuitive hierarchy for auditability of common multi-unit auction formats. We perform a brief empirical exercise to support our theoretical results. Using the classification of government debt sales from Brenner et al. [2009] and the 2017 Corruption Perception Index, we show that corruption is positively associated with credibility.[7] One might expect that more corrupt nations should implement less credible mechanisms, to better extract the rents of corruption. However, corruption is measured broadly across an economy, not just within treasury auctions, and perceived corruption disincentivizes market participation.[8] To successfully raise funds, nations perceived as corrupt should implement credible mechanisms for selling sovereign debt. To the extent that corruption is a feature of a society and not just its government, the treasury has a vested interest in making bidder collusion more difficult. Releasing less post-auction information, possible in a less auditable auction format, inhibits the ability of bidders to verify each others' actions.[9] This possibility suggests that corruption should be positively correlated with auditability, as in the case of providing a guarantee against government corruption. Our empirical analysis supports these hypotheses.

Although our results are constrained to the analysis of static, deterministic multi-unit auctions, our auditability hierarchy is generally applicable. As discussed above, the question of how much information must be released is a natural one, applicable to many other contexts. Contextual details may effect the strength of the results — the sharpness of our ordering in some cases depends on the ability of the auctioneer to arbitrarily increase supply — but we expect the rough intuition that it requires less information to verify outcomes that depend only on own action than those that depend on an action profile to hold generally.

[7] This directionally replicates the results of Brenner et al. [2009], who do not find statistical significance.

[8] As examined in [McAdams and Schwarz, 2007b], all "non-belief" costs are borne by the auctioneer when bidders do not believe he will adhere to the stated rules.

[9] To prevent the release of information valuable for collusion, the auctioneer could apply differential privacy techniques [McSherry and Talwar, 2007]. Uncertainty is not a feature in our model since conditional on public information a deterministic mechanism is either believable or not.

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 known about 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, 2018, 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 "practical" model of bidder values that provides an identical revenue comparison across observed implementations. We contribute to this literature by showing that, although revenue may not be comparable, the discriminatory auction is easier to auditably implement than a uniform price auction. In some circumstances, this may constrain practical mechanism selection (see our discussion of corruption in the introduction). 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, 2018, 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 public information release, 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 is related insofar as it 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. LiCalzi and Pavan [2005] show that in uniform price auctions an elastic supply curve can improve revenue by inducing competition at low prices. Committing to supply

that increases in market clearing price shrinks the set of equilibrium market clearing prices, and can eliminate severe equilibrium underpricing. McAdams [2007] considers a uniform price auctioneer who selects a quantity after bids are submitted, and prior to the auction buyers know only the auctioneer's (true) cost curve. In this case the market clearing price is almost always equal to bidders' value for the goods. We focus instead on the role that potentially infinite supply has in bidders determining if the auction was run as planned, and ignore equilibrium bidding.[10]

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. As mentioned above, we focus on comparability of information release independent of bidder equilibrium, so there is not a direct tie to these results. Akbarpour and Li [2018] define credibility, where a mechanism is credible if the auctioneer cannot improve his own outcome without alerting the agents. Our notion of auditability provides a hierarchy of mechanisms, which augments the binary notion of credibility; the previous note on equilibrium inference also applies.[11]

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 while our information may be value-relevant, we do not analyze it in these terms). Our auctioneer's evaluation method is more basic: can he improve his outcome without

[10] Our results continue to hold when the auctioneer announces an elastic supply schedule rather than a specific quantity. We assume fixed inelastic supply to simplify our analysis.

[11] 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 subsequent to the mechanism being run. The latter is a natural interpretation of "daylight" provisions regarding transparency of government processes.

tipping off bidders to his deviation, independent of the information he learns from equilibrium bidding?[12]

Finally, our analysis of incentive-compatible auditing ties to work on finding mechanisms that minimize (valid) participant complaints. Abdulkadiroğlu and Sönmez [2003] demonstrate 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 τ-fairness, relating to the ability of rematches implied by justified envy to themselves generate justified envy. Troyan et al. [2018] require 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.

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 $\theta_i \in \Theta_i$. θ_i is identically distributed for all bidders, and θ_i is independent of θ_j for all bidders $i \neq j$. 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 quantities.

After observing θ_i each bidder i submits bid $b_i \in B$, yielding bid profile $b = (b_i, b_{-i})$, and (personal) outcome $o^i(b) = (q^i(b), t^i(b))$ is realized. 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

[12] McAdams and Schwarz [2007b] show that the cost of non-believability is borne entirely by the seller; this result is echoed by McAdams and Schwarz [2007a], who show that it is possible for a reputable intermediary to improve sales efficiency. This suggests that auctioneers have a vested interest in their mechanisms being auditable. If information release is generally undesirable, more auditable mechanisms will be preferred to less auditable mechanisms.

actions fixed across auction formats, we refer to o as a *mechanism.*[13] 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 $u^0(b) = \sum_{i=1}^n t^i(b_i, b_{-i})$.

After outcomes are realized, the auctioneer releases public information $I = \mathcal{I}(b)$, where $\mathcal{I} : B^n \rightrightarrows B^n$. We refer to $\mathcal{I}$ as a public information release. We assume that public information is consistent with the bid profile, $b \in \mathcal{I}(b)$, and that $\mathcal{I}$ yields 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 bid profile b, if there is some $b'_{-i} \in B^{n-1}$ such that $o'_i = o^i(b_i, b'_{-i})$, and $(b_i, b'_{-i}) \in \mathcal{I}(b)$. Outcome profile o' is *explicable*, given bid profile b, if o'_i is explicable for each bidder i.

Denote the set of outcomes explicable for bidder i by $X^i(b)$, and the set of explicable outcomes by $X(b) = \times_{i=1}^n X^i(b_i)$. An outcome o' is explicable given bid profile b if it is explicable for each bidder, but it may not be the case that there is $b' \in \mathcal{I}(b)$ such that $o' = o(b')$. Note that since $b \in \mathcal{I}(b)$ and $o(b)$ is explicable for each bidder i, $X(b) \supseteq \{o(b)\}$ is nonempty. Explicability determines the set of deviations from plan the seller could implement, conditional on a particular bid profile.

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

$$o\left(b\right) \in \arg\max_{o' \in X(b)} u^0\left(o'\right).$$

That is, the auction is audited by public information release $\mathcal{I}$ if the outcomes generated by the auction maximize the auctioneer's utility, subject to consistency with each bidder's knowledge of her own action. 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.

[13] 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.

We do not permit the auctioneer to misreport public information. This is a natural restriction when public information regards sufficient statistics for market outcomes, such as quantities and prices.

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

Remark 1. 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.

The auction game being audited by public information release is distinct from credibility of the meta-game in which the auction is augmented by public information. Credibility (in the sense of Akbarpour and Li [2018]) presumes that the auctioneer is best 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 [2018] and our introduction, equilibrium bids take very different forms in different multi-unit auction formats, so requiring auditability of only on-path bids would void the ability to compare the auditability of different auction formats.

Although public information is mathematically structured as a partition on the space of bid profiles, in practice information tends to more related to measured observables. 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. In the case of deterministic mechanisms focusing on a partition of the bid space is sufficient for this kind of information release, since outcomes are identified with bid profiles, and enables the comparison of information release across mechanisms. Conditional on a particular bid, certain outcomes may be feasible in one auction that are not feasible in another, so focusing only on outcomes would not permit comparison across mechanisms. Additionally, there are deterministic mechanisms that cannot be audited with information about outcomes, even allowing for inference from one's own submitted bid.[14] Constraining attention to deterministic mechanisms,

[14] 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

the analysis of information release in terms of bid partitions is therefore necessary and sufficient. With nondeterministic mechanisms explicit information about outcomes, or the process generating random outcomes, may be necessary.

Definition 3. [More auditable] Auction o is *more auditable* than auction o' if for any public information release $\mathcal{I}$ that audits o', $\mathcal{I}$ audits o. We write $o \trianglerighteq o'$. If $o \trianglerighteq o'$ and there is a public information release $\mathcal{I}$ that audits o and not o', then auction o is *strictly more auditable* than auction o', and we write $o \triangleright o'$.

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 ordering of auditability to show that menu mechanisms are more auditable than any other mechanism, and are the revenue-maximizing mechanisms auditable in zero information. For comparison across mechanisms our auction model assumes identical bid spaces, independent of the outcome function. We therefore have no innate ability to distinguish between 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 a menu mechanism.

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 public information release $\mathcal{I}$.

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

own bid.

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 $o'_i = o^i(b_i, b'_{-i})$. 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 the set of explicable outcomes $X(b)$ has $o'_i = o^i(b)$ for all bidders i, and hence $X(b) = \{o(b)\}$. Then $o(b) \in \arg\max_{o' \in \{o(b)\}} u^0(o')$. □

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

Similar arguments show that any mechanism auditable in zero information is revenue-equivalent to a menu mechanism.

Theorem 2. [Zero-auditable mechanisms are menus] *Let o be a mechanism. If $\mathcal{I}$ audits o for any information release function $\mathcal{I}$, 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

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

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}. □

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. As explored in Akbarpour and Li [2018], first price auctions are creditable, however in our model they are not auditable in zero information. The analysis in Akbarpour and Li [2018] crucially assumes that there is a known maximum supply available for auction. We view this as a natural target for auditing. As discussed in the introduction, 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.[15] Consider seller who claims to be selling a single unit in a first-price auction. If bidders are mistaken

[15] For example, in a first price auction between two bidders, each of whom demands one unit, if two units are available both bids are zero.

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.[16]

3.1 Common multi-unit auctions

We begin by defining 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 bid vectors $b \in \mathbb{R}^Q$, and awards units to the Q highest bids. When tiebreaking is necessary, we assume a deterministic tiebreaking rule is used, independent of the auction format. Given supply Q and bid profile b, the last accepted bid and first rejected bid are

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

Intuitively, the last accepted bid is the lowest market-clearing price such that the market has weak excess supply and the first rejected bid is the highest market-clearing price such that the market has strict excess demand.[17] Accordingly, we refer to the highest and lowest market-clearing prices $p^\star_{\text{LAB}}$ and $p^\star_{\text{FRB}}$, and a price $p^\star$ is a *market-clearing price* if $p^\star_{\text{LAB}} \geq p^\star \geq p^\star_{\text{FRB}}$. In standard multi-unit auctions all bids $b_{iq} > p^\star_{\text{FRB}}$ are awarded, and if $p^\star_{\text{LAB}} > p^\star_{\text{FRB}}$ then tiebreaking is unnecessary.

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

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

In a uniform price auction, bidders pay a market-clearing price for each unit they receive. Following the analysis in Burkett and Woodward [2018] we consider the last accepted bid and first rejected

[16] Akbarpour and Li [2018] show that this format can be credible if the seller is truly and publicly quantity-constrained. 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.

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

bid uniform price auctions, where

$$o^i_{\text{LAB}}(b) = \left(q^i(b), t^i_{\text{LAB}}(b)\right), \qquad t^i_{\text{LAB}}(b) = p^\star_{\text{LAB}}(b)\, q^i(b);$$
$$o^i_{\text{FRB}}(b) = \left(q^i(b), t^i_{\text{FRB}}(b)\right), \qquad t^i_{\text{FRB}}(b) = p^\star_{\text{FRB}}(b)\, q^i(b).$$

To intuitively compare the auditability of different auction formats we consider the inferences necessary to audit a mechanism. 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 $o(b) \in \arg\max_{o' \in X(b)} u^0(o')$. Then if $\mathcal{I}$ audits o_{PAB}, $\mathcal{I}$ need only make explicit aggregate allocations, and then only in certain circumstances. Corner cases notwithstanding, announcing quantity allocations is sufficient to audit the discriminatory auction.

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}.*

Proof. Suppose that o'_i is explicable for agent i, given bid profile 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 $t'_i = \sum_{q=1}^{q'_i} b_{iq} = t^i(b)$. Then $o'_i = o^i(b)$ and there is a unique explicable outcome, $X(b) = \times_{i=1}^n \{o^i(b)\} = \{o(b)\}$. Then $o(b) \in \arg\max_{o' \in \{o(b)\}} u^0(o')$, and $\mathcal{I}$ audits o_{PAB}. □

There is a natural sense in which 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 is also necessary, 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.

Theorem 3. [PAB more auditable than LAB] *The discriminatory auction is more auditable than the last accepted bid uniform price auction, $o_{PAB} \trianglerighteq o_{LAB}$. If $Q > 1$, this relation is strict.*

Proof. We first show that if $\mathcal{I}$ audits o_{LAB}, then $\mathcal{I}$ audits o_{PAB}. Let b be a bid profile such that $o_{\text{LAB}}(b) \in \arg\max_{o' \in X_{\text{LAB}}(b)} u^0(o')$ but $o_{\text{PAB}}(b) \notin \arg\max_{o' \in X_{\text{PAB}}(b)} u^0(o')$. Because, conditional on bids, the discriminatory auction maximizes revenue under the constraint $\sum_{i=1}^n q_i = Q$, it must be that there is $o' \in X_{\text{PAB}}(b)$ with $\sum_{i=1}^n q'_i > Q$ and $b_{iq_i} > 0$ for each bidder i. Then there is some bidder i for whom $q'_i > q^i(b)$, and therefore there is an opponent bid profile b'_{-i}, $(b_i, b'_{-i}) \in \mathcal{I}(b)$, such that $q^i(b_i, b'_{-i}) = q'_i > q^i(b)$. Now consider the bid profile $b' = (b_i, b'_{-i})$. It must be that $p^\star_{\text{LAB}}(b') \leq p^\star_{\text{LAB}}(b)$, since $b' \in \mathcal{I}(b)$ and $\mathcal{I}$ audits o_{LAB}. This inequality cannot be strict, since $b \in \mathcal{I}(b')$ and $\mathcal{I}$ audits o_{LAB}. Then $p^\star_{\text{LAB}}(b') = p^\star_{\text{LAB}}(b)$. But note that q' is explicable, independent of the auction mechanism, since the quantity allocation rule is identical in both mechanisms. Then since $q'_i > q^i(b)$, it follows that $\tilde{o}'$, where $\tilde{o}'_i = (q'_i, q'_i p^\star_{\text{LAB}}(b))$ and $\tilde{o}'_j = (q^j(b), t^j(b))$ for all $j \neq i$, is explicable given bid profile b. This improves the seller's revenue, contradicting the assumption that $\mathcal{I}$ audits o_{LAB}. Then $o_{\text{PAB}} \trianglerighteq o_{\text{LAB}}$.

We now show that if $Q > 1$ there is an information release $\mathcal{I}$ that audits o_{PAB} but not o_{LAB}. Let $\mathcal{I}(b)$ be the set of bid functions that generate allocation $q(b)$ when the auction is run according to the stated rules. Lemma 1 implies that $\mathcal{I}$ audits o_{PAB}. Consider bidders i and j, $i \neq j$, and let $q_i, q_j > 0$ be such that $q_i + q_j = Q$. Let $b(p_i, p_j) = (b_i(p_i), b_j(p_j), b_{-ij})$ be such that

$$b_{iq}(p_i) = \begin{cases} p_i & \text{if } q \leq q_i \\ 0 & \text{if } q > q_i; \end{cases} \quad b_{jq}(p_j) = \begin{cases} p_j & \text{if } q \leq q_j, \\ 0 & \text{if } q > q_j; \end{cases} \quad b_{kq} = 0 \ \ (k \neq i, j).$$

By construction, $b(p'_i, p'_j) \in \mathcal{I}(b(p_i, p_j))$ for all p, p' such that $b_i(p_i), b_j(p_j) \in B$. For any p, $p^\star_{\text{LAB}} = \min\{p_i, p_j\}$. Assume without loss of generality that $p_i > p_j$. Then $b(p_i, p_i) \in \mathcal{I}(b(p_i, p_j))$, and $(q_i, p_i q_i)$ is explicable given $b(p_i, p_j)$. This generates strictly greater revenue than $(q_i, p^\star_{\text{LAB}} q_i)$, hence $\mathcal{I}$ does not audit o_{LAB}. □

The quantity restriction in Theorem 3 is essential. Mathematically, the first step of the proof fails because we cannot assume $q_i, q_j > 0$ for two agents $i \neq j$. Economically, when a single unit is for sale there is no distinction between a discriminatory auction and a last accepted bid uniform price auction, since the last (and only) accepted bid is the first bid.

While the intuition, in terms of what needs to be revealed, for the discriminatory auction being strictly more auditable than the last accepted bid uniform price auction is fairly straightforward,

the proof of Theorem 3 takes a slightly different approach. This is due to a number of corner cases in the "quantity must be verified" argument. In particular, public information release can audit a discriminatory auction even without verifying the quantity sold. Suppose that if there is excess demand the entire bid profile is made public, otherwise it is simply publicized that there is weak excess supply. Then in the latter case each bidder knows only her own quantity allocation and not those of her opponents, but the auction outcome still maximizes the seller's revenue in the space of explicable outcomes.

Theorem 4. [PAB strictly more auditable than FRB] *The discriminatory auction is strictly more auditable than the first rejected bid uniform price auction, $o_{PAB} \rhd o_{FRB}$.*

Proof. The proof that $o_{\text{PAB}} \trianglerighteq o_{\text{FRB}}$ is substantially similar to the proof of Theorem 3 ($o_{\text{PAB}} \trianglerighteq o_{\text{LAB}}$) and is omitted here. Now let $\mathcal{I}$ be such that $\mathcal{I}(b) = \{b' : \ q(b') = q(b)\}$. By Theorem 1, $\mathcal{I}$ audits o_{PAB}. Define the bid profile $b(p, p') = (b_i(p), b_{-i}(p'))$ by

$$b_{jq}(p) = \begin{cases} p & \text{if } q \leq Q, \\ 0 & \text{if } q > Q, \end{cases} \quad \text{for all } j \in \{1, \ldots, n\}.$$

Note that for all $p > p'$, $q^i(b(p, p')) = Q$ and $q^j(b(p, p')) = 0$ for all $j \neq i$. Then $b(\tilde{p}, \tilde{p}') \in \mathcal{I}(b(p, p'))$ whenever $\tilde{p} > \tilde{p}'$ and $p > p'$. Then for all $p > p'$, $o^i(b(p, p')) = (Q, p'Q)$. It follows that $o(b(p, 0)) \notin \arg\max_{o' \in X_{\text{FRB}}(b(p,0))} u^0(o')$, and $\mathcal{I}$ does not audit o_{FRB}. □

Surprisingly, without further restrictions the last accepted bid auction is not more auditable than the first rejected bid auction. In both uniform price auctions we analyze 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, 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 (or set of bidders). If $p^\star_{\text{LAB}}(b) > p^\star_{\text{FRB}}(b)$ the auctioneer can increase the market clearing price for all bidders, not just idiosyncratically, and increase his revenue.[18] It would thus seem that the last accepted bid auction is more auditable than

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

a first rejected bid auction. However, 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 and FRB cannot be ranked] *If $Q > 1$, the last accepted bid and first rejected bid uniform price auctions cannot be ranked in auditability.*

Proof. We give an information release that audits o_{FRB} but not o_{LAB}. Let $\mathcal{I}_{\text{FRB}}$ give the set of bids which yield a particular quantity allocation and first rejected bid,

$$\mathcal{I}_{\text{FRB}}(b) = \left\{b' : q(b') = q(b) \text{ and } p^{\star}_{\text{FRB}}(b') = p^{\star}_{\text{FRB}}(b)\right\}.$$

By construction, $\mathcal{I}_{\text{FRB}}$ audits o_{FRB}. Let $q_i, q_j > 0$ be such that $q_i + q_j = Q$, and consider the bid profile b from the proof of Theorem 3. Then $p^{\star}_{\text{FRB}}(b') = 0$ for all $b' \in \mathcal{I}_{\text{FRB}}(b)$. However, as in the proof of Theorem 3 there is an explicable outcome $o' \in X(b)$ where $o'_i = (q_i, p_i q_i)$ and $o'_j = (q_j, p_j q_j)$. When $p_i > p_j$, $p^{\star}_{\text{LAB}}(b(p_i, p_j)) = p_j$, and o' yields greater seller revenue than $o(b)$. Then $\mathcal{I}_{\text{FRB}}$ does not audit o_{FRB}.

A similar construction gives an information release $\mathcal{I}_{\text{LAB}}$ that audits o_{LAB} but not o_{FRB}. □

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

Proposition 1. [Hierarchy of auditability] *Let o_{MENU} be any menu mechanism. Then $o_{MENU} \rhd o_{PAB} \trianglerighteq o_{LAB}, o_{FRB}$. If $Q > 1$, the weak relation is strict.*

The space of possible multi-unit auction mechanisms is large, and Proposition 1 is an incomplete characterization. Nonetheless it orders the most commonly observed multi-unit auction formats. It can also be shown that Spanish and Vickrey auctions are less easily auditable than discriminatory auctions, but cannot themselves be ranked. Auditing a Spanish auction requires the ability to compute the mean winning bid, while auditing a Vickrey auction requires knowledge of losing bids. We leave a more thorough categorization of auditability to future work.

3.2 Constrained auditability

There is a large number of partitions of the bid space, many of which lack a clear economic motivation. We introduce the notion of constrained auditability, where one mechanism is more auditable

than another if it is audited by any information release, from within a particular class, which audits the other. With constrained auditability, we focus on information release functions with particular economic intuitions.

Definition 6. [Constrained more auditable] Let P be a set of information release functions. Mechanism o is *P-constrained more auditable* if any $\mathcal{I} \in P$ that audits o' also audits o, and we write $o \trianglerighteq_{P\emptyset} o'$.

Lemma 2. [Heritability of hierarchy] *Let P be a set of information release functions, and suppose that $o \trianglerighteq_{P\emptyset} o'$. Then for any $P' \subseteq P$, $o \trianglerighteq_{P'\emptyset} o'$.*

Relative auditability is inherited down public information constraints because it is a property of all information release functions in a given set. Constraining the set, relative auditability will still hold. Strict relative auditability is not inherited, since it is defined from the existence of a public information release with a particular property, and shrinking the set of feasible information release functions may remove those with this property.

In practice many auctioneers release information about the market clearing price $p^\star_{\text{LAB}}$. Recall that in each of the common auction formats above, all bids strictly above $p^\star_{\text{LAB}}$ are accepted and all bids strictly below $p^\star_{\text{LAB}}$ are rejected. Conditional on releasing this information, discriminatory auctions and last accepted bid uniform price auctions are equally auditable.

Theorem 6. [Hierarchy conditional on market price] *Let P be the set of information release functions $\mathcal{I}$ such that for any bid profile b, $\mathcal{I}(b) \subseteq \{b' : p^\star_{LAB}(b') = p^\star_{LAB}(b')\}$. Then if o_{MENU} is any menu mechanism,*

$$o_{MENU} \rhd_{P\emptyset} o_{PAB} \sim_{P\emptyset} o_{LAB} \rhd_{P\emptyset} o_{FRB}.$$

Proof. Lemma 2 implies that we need only show that $o_{\text{LAB}} \trianglerighteq_{P\emptyset} o_{\text{PAB}}$, that there is $\mathcal{I} \in P$ that audits o_{PAB} and not o_{FRB}, and that there is $\mathcal{I} \in P$ that audits o_{MENU} and not o_{PAB}.

Suppose that $\mathcal{I} \in P$ audits o_{PAB}. If $\mathcal{I}$ does not audit o_{LAB}, there is a bid profile b such that $o(b) \notin \arg\max_{o' \in X_{\text{LAB}}(b)} u^0(o')$. Because $p^\star_{\text{LAB}}(b') = p^\star_{\text{LAB}}(b)$ for all $b' \in \mathcal{I}(b)$, explicability implies that $o'_i = (q'_i, p^\star_{\text{LAB}}(b) q'_i)$ for some q'_i. Since $o(b)$ does not maximize the seller's revenue, it must be that $q'_i > q^i(b)$ for some bidder i, and $p^\star_{\text{LAB}}(b) > 0$; since o' is explicable it must be that $b_{iq'_i} = p^\star_{\text{LAB}}(b)$. But since the quantity allocation $q(\cdot)$ does not depend on auction format, there is

an explicable outcome $\tilde{o}' \in X_{\text{PAB}}(b)$ with $\tilde{q}'_i = q'_i > q^i(b)$. With $b_{iq'_i} > p^\star_{\text{LAB}}(b)$, it follows that $o(b) \notin \arg\max_{o' \in X_{\text{PAB}}(b)} u^0(o')$, contradicting the assumption that $\mathcal{I}$ audits o_{PAB}. Then $o_{\text{LAB}} \trianglerighteq_{P\emptyset} o_{\text{PAB}}$, and Lemma 2 and Theorem 3 imply that $o_{\text{PAB}} \sim_{P\emptyset} o_{\text{LAB}}$.

Second, let $\mathcal{I}(b) = \{b' : \ p^\star_{\text{LAB}}(b') = p^\star_{\text{LAB}}(b) \text{ and } q(b') = q(b)\}$. By Lemma 1, $\mathcal{I}$ audits o_{PAB}. Consider the bid profile $b(p) = (b_i(p), b_{-i})$, where

$$b_{iq}(p) = p, \ \ b_{jq} = 0 \ \ (j \neq i).$$

Note that $p^\star_{\text{LAB}}(b(p)) = p$. Quantity allocations and $p^\star_{\text{LAB}}$ do not change if $p > b'_{-i} = p' > 0$, so $o'_i = (Q, p'Q)$ is explicable for bidder i. Since $p'Q > 0$, $\mathcal{I}$ does not audit o_{FRB}.

Finally, let $\mathcal{I}(b) = \{b' : \ p^\star_{\text{LAB}}(b') = p^\star_{\text{LAB}}(b)\}$. By Theorem 1, $\mathcal{I}$ audits o_{MENU}. Consider the bid profile $b(p) = (b_i(p), b_j(p), b_{-ij})$, where

$$b_{iq}(p) = p, \ \ b_{jq}(p) = p, \ \ b_{kq} = 0 \ \ (k \neq i, j).$$

Note that $p^\star_{\text{LAB}}(b(p)) = p$. Because this is the market-clearing price if bidder i (or j) submits her bid while all other bidders bid 0, $o'_\ell = (Q, pQ)$ is explicable for bidders $\ell \in \{i, j\}$. Since $2pQ > pQ$, $\mathcal{I}$ does not audit o_{PAB}. □

3.3 Incentive compatible auditability

An outcome is explicable if it is consistent with public information and each bidder's own action. Public information release audits a mechanism if the "correct" outcome is revenue maximizing in the space of explicable outcomes. This implicitly assumes that inexplicable outcomes are somehow costly. This assumption is motivated by the idea that if a bidder sees something inexplicable, they will report it to a central authority, perhaps because suspicion in one outcome may generate suspicions of other outcomes (that were otherwise explicable). However, it is also possible that bidders will complain only if inexplicable outcomes generate less utility than explicable outcomes.

Definition 7. [Pleasant surprises] Given bid profile b, bidder i is *pleasantly surprised* by outcome o'_i if $\bar{u}^i(o'; \theta_i) \geq \bar{u}^i(\tilde{o}'; \theta_i)$ for all explicable outcomes $\tilde{o}'_i \in X^i(b)$ and almost all types θ_i.

A bidder is pleasantly surprised by an outcome if it generates at least her best explicable utility,

regardless of whether it is explicable. Implicit in the definition of auditability is the assumption that there is something wrong with outcomes being inexplicable. If a bidder witnesses an outcome she cannot explain, she knows that the designer's claimed mechanism was not run. We do not model what happens if the bidder learns that the designer's claims were not honored, but it is reasonable to suspect that this will cause problems for the designer, perhaps because the bidder reports the deviation to a central authority. However, we might alternately assume that bidders report deviations only if their inexplicable outcome is *bad*: a pleasant surprise generates more utility than the bidder thought possible, so there is no direct incentive to complain.

Definition 8. [No complaints] Given bid profile b, outcome o_i' *generates no complaints from bidder* i if it is explicable or a pleasant surprise. Given bid profile b, outcome o *generates no complaints* if it generates no complaints from any bidder i.

Denote the set of outcomes which generate no complaints from bidder i by $Y^i(b)$, and the set of outcomes which generate no complaints (from any bidder) by $Y(b) = \times_{i=1}^n Y^i(b)$.

Definition 9. [Auditable without complaints] The mechanism o is *audited without complaints* by public information release $\mathcal{I}$ if for all bid profiles b,

$$o(b) \in \arg\max_{o' \in Y(b)} u^0\left(o'\right).$$

If any public information release $\mathcal{I}$ that audits o' without complaints also audits o without complaints, we write $o \trianglerighteq^{\text{IC}} o'$.

Whether a mechanism is auditable without complaints depends on the richness of the agents' utility space. When utility is relatively known (the space of utility functions is limited) the auctioneer may be relatively certain that a particular inexplicable outcome leaves auction participants no worse off. For example, suppose that bidders demand at most one unit, and each bidder has value uniformly distributed on $[1, 2]$. By Lemma 1, a first-price auction for a single unit can be audited by announcing the resulting quantity allocation. The auctioneer, however, can make all losing bidders weakly better off by giving a good to each at price 1. This is a pleasant surprise for each losing bidder, the winner's outcome is explicable, and the seller's revenue improves by $n - 1$. It follows that in this utility space announcing quantities does not audit without complaints the first-price

auction. The key assumption is that an outcome generates no complaints if it is explicable or a pleasant surprise for *almost all* types.

Theorem 7. [Auditable without complaints is auditable] *Let $\bar{v} \in \mathbb{R}_{++}$ and suppose that for any $v : \mathbb{N} \times \mathbb{R} \to (-\infty, \bar{v})$ that is weakly increasing in quantity and strictly decreasing in transfer there is $\theta \in \Theta$ such that $\bar{u}(\cdot, \cdot; \theta) \equiv v$. Then mechanism o is audited without complaints by public information release $\mathcal{I}$ if and only if o is audited by $\mathcal{I}$.*

Proof. By definition, $Y(b) \supseteq X(b)$ for any bid profile b. Then regardless of the underlying utility space, if $\mathcal{I}$ audits o without complaints, $\mathcal{I}$ audits o.

Now suppose that $\mathcal{I}$ audits o. If $\mathcal{I}$ does not audit o without complaints, there is a bid profile b such that $o(b) \notin \arg\max_{o' \in Y(b)} u^0(o')$; then since $\mathcal{I}$ audits o, there is an outcome o' such that $u^0(o') > u^0(o(b))$, hence a bidder i such that $t_i' > t^i(b)$, where $o_i' \notin X^i(b)$ is a pleasant surprise for bidder i. Because utility is decreasing in transfer and increasing in quantity, it must be that $q_i' > q^i(b)$. Because the utility space is rich, $\Pr(\bar{u}^i(q^i(b), t^i(b); \theta_i) = \bar{u}^i(q_i', t^i(b); \theta_i)) > 0$, implying that $\bar{u}^i(o^i(b); \theta_i) > \bar{u}^i(o_i'; \theta_i)$ with positive probability. Then o_i' is not a pleasant surprise, contradicting the assumption that o_i' is a pleasant surprise. Then $\mathcal{I}$ audits o without complaints. □

When the type space is rich, the auctioneer cannot know, given a particular bid, that a particular outcome improves a bidder's utility. Then bidders might report an unexpected outcome to a central authority. With a restricted type space this may no longer be feasible. For example, if bidders have strictly positive marginal utility above $m > 0$ for each additional unit received, the auctioneer can issue additional units at will at a price of $m/2$ to obtain additional profits. In equilibrium the seller may be able to infer information about bidder values, and thereby avoid this problem. Because we do not analyze bidder equilibrium behavior, we preclude this possibility in our analysis.[19] On the other hand, the richness requirement of Theorem 7 is over-strong; the result holds equally well in an unbounded one-dimensional type space where bidders have maximum demand given by θ_i.

[19] Allowing inference from equilibrium strategies invalidates the equivalence, and indeed makes auditability quite difficult. Consider a first-price auction between two bidders. If the auctioneer surprises the loser with a unit at half her bid, the loser is pleasantly surprised and the auctioneer is strictly better off in almost all cases. Disallowing inference from bids means that while the auctioneer can charge half the losing bid, he cannot be certain (as he would be in equilibrium) that the loser's value is below half this bid, and therefore this inexplicable allocation may generate a complaint.

3.4 Outcome-constrainted auditability

The sharpness of our results — for example, Theorem 2's implication that menu mechanisms are the unique mechanisms auditable without public information — relies on the assumption that the feasible outcome space is $O = (\mathbb{N} \times \mathbb{R})^n$. As discussed in the introduction, this is a natural outcome space under the assumption that supply may be increased without bound, as when goods are replicable. In practical single-unit auctions, however, it is frequently clear that there is a single unit available (a single painting, a single construction contract, etc.). We therefore consider the notion of auditability when participants have some foreknowledge of the feasible outcome space.

Definition 10. [Auditing under outcomes] The mechanism o is *audited under outcomes* $\tilde{O}$ by public information release $\mathcal{I}$ if for all bid profiles b,

$$o(b) \in \arg \max_{o' \in X(b) \cap \tilde{O}} u^0(o').$$

If any public information release $\mathcal{I}$ that audits o' under outcomes $\tilde{O}$ also audits o under outcomes $\tilde{O}$, we write $o \trianglerighteq_{\emptyset \tilde{O}} o'$.

Note that o is audited under outcomes $\tilde{O}$ (given any public information release) only if $o(b) \in \tilde{O}$. Because $o' \in \tilde{O}$ constrains the search space for the seller's maximization problem, it is immediate that auditability is essentially heritable.

Lemma 3. [Heritability of hierarchy'] *Suppose that $o \trianglerighteq o'$. Then if $o(b), o'(b) \in \tilde{O}$ for all bid profiles b, $o \trianglerighteq_{\emptyset \tilde{O}} o'$.*

Proof. Note that, for any mechanism $\tilde{o}$, if $\mathcal{I}$ audits $\tilde{o}$ and $o(b) \in \tilde{O}$ for all bid profiles b, then $\mathcal{I}$ audits $\tilde{o}$ under outcomes $\tilde{O}$. The result is then immediate. □

Theorem 8. [Hierarchy with supply commitment] *Fix Q, and let $\tilde{O} = \{(q_i, t_i)_{i=1}^n : \sum_{i=1}^n q_i \leq Q\}$. For any menu mechanism o_{MENU},*

$$o_{MENU} \sim_{\emptyset \tilde{O}} o_{PAB} \trianglerighteq_{\emptyset \tilde{O}} o_{LAB}, \text{ and } o_{PAB} \triangleright_{\emptyset \tilde{O}} o_{FRB}.$$

If $Q > 1$, $o_{PAB} \triangleright_{\emptyset \tilde{O}} o_{LAB}$.

Proof. In light of Lemma 3, we need only show that o_{PAB} is auditable under outcomes $\tilde{O}$ in zero information, that there is public information release $\mathcal{I}$ that audits o_{PAB} but not o_{FRB} under outcomes $\tilde{O}$, and that when $Q > 1$ there is public information release $\mathcal{I}$ that audits o_{PAB} but not o_{LAB} under outcomes $\tilde{O}$.

First, recall that the discriminatory auction for Q units accepts the Q highest bids. Since the seller has committed to sell no more than quantity Q (via $\tilde{O}$), the only way to strictly improve revenue is to allocate additional units (if less than Q) is sold, or to accept a higher bid that was previously rejected in favor of a lower one. The latter possibility is not possible in the discriminatory framework, and less than Q cannot be sold. It follows that $o_{\text{PAB}}(b)$ maximizes $u^0(o')$ on $o' \in \tilde{O}$. Then o_{PAB} is zero-auditable.

Second, let $\mathcal{I}$ be such that $\mathcal{I}(b) = B^n$ for all b. The proof that $o_{\text{PAB}} \rhd o_{\text{LAB}}$ (Theorem 3) is equally valid in this case. Then $\mathcal{I}$ does not audit o_{LAB} without complaints under outcomes $\tilde{O}$, and $o_{\text{PAB}} \rhd_{\emptyset\tilde{O}} o_{\text{LAB}}$.

Finally, o_{FRB} is not zero-auditable, and the construction Theorem 6 is equally valid in this case. Then $o_{\text{LAB}} \rhd_{\emptyset\tilde{O}} o_{\text{FRB}}$. □

Outcome-constrained auditability is in some sense the closest auditability comes to credibility (Akbarpour and Li [2018]), and in this light Theorem 8 is analogous to the claim that discriminatory auctions are credible, and uniform price auctions are not. However auditability is a stronger constraint than credibility, since a mechanism must be explicably revenue-maximizing for all bid profiles, not just those that arise in equilibrium.

Each of our restrictions on auditability is individually instructive, but taken in concert they provide a natural ranking on the information necessary to believe that an auction was run honestly.

Definition 11. The mechanism o is *audited without complaints under outcomes* $\tilde{O}$ by public information release $\mathcal{I}$ if for all bid profiles b,

$$o(b) \in \arg \max_{o' \in Y(b) \cap \tilde{O}} u^0\left(o'\right).$$

Given a set of information release functions P, mechanism o *P-audit dominates* mechanism o' if any public information release $\mathcal{I} \in P$ that audits o' without complaints under outcomes $\tilde{O}$ also audits o

without complaints under outcomes $\tilde{O}$, $o \trianglerighteq^{\text{IC}}_{P\tilde{O}} o'$.

Proposition 2. [Practical auditability hierarchy] *Suppose that the utility space is rich in the sense of Theorem 7. Fix Q and let $\tilde{O} = \{(q_i, t_i)_{i=1}^n : \sum_{i=1}^n q_i \leq Q\}$, and let P be the set of information release functions $\mathcal{I}$ such that $\mathcal{I}(b) \subseteq \{b' : p^\star_{LAB}(b') = p^\star_{LAB}(b)\}$. Then for any menu mechanism o_{MENU},*

$$o_{MENU} \sim^{IC}_{P\tilde{O}} o_{PAB} \sim^{IC}_{P\tilde{O}} o_{LAB} \rhd^{IC}_{P\tilde{O}} o_{FRB}.$$

If $P' = P \cup \{B^n\}$ and $Q > 1$,

$$o_{MENU} \sim^{IC}_{P'\tilde{O}} o_{PAB} \rhd^{IC}_{P'\tilde{O}} o_{LAB} \rhd^{IC}_{P'\tilde{O}} o_{FRB}.$$

Proof. This result follows from the same methods applied in Theorems 6, 7, and 8. and is omitted here. □

We also point out that the preceding arguments are not meaningfully affected by the presence of a reserve price.[20]

Corollary 3. [Hierarchy of supply commitment with reserve] *Proposition 2 holds in the presence of any reserve price $r < \bar{v}$.*

Proposition 2 implies that in (strict) multi-unit auctions, where the quantity is committable, under common public information structures the discriminatory auction is more auditable than common uniform price auctions, and within the class of uniform price auctions the last accepted bid auction is more auditable than the first rejected bid auction. Provided the space of possible bidder utility functions is sufficiently rich, this statement holds even if inexplicability is a problem only when outcomes are potentially utility-disimproving. Thus in practice we might expect measures positively correlated with trustworthiness to be negatively correlated with implementation of a discriminatory auction; controlling for implementation of a uniform price auction, measures positively correlated with trustworthiness should be negatively correlated with a last accepted bid pricing rule.

[20] The precise constructions in the proof sometimes rely on the ability to bid 0 for certain units. This can be replaced with bidding the reserve price without issue.

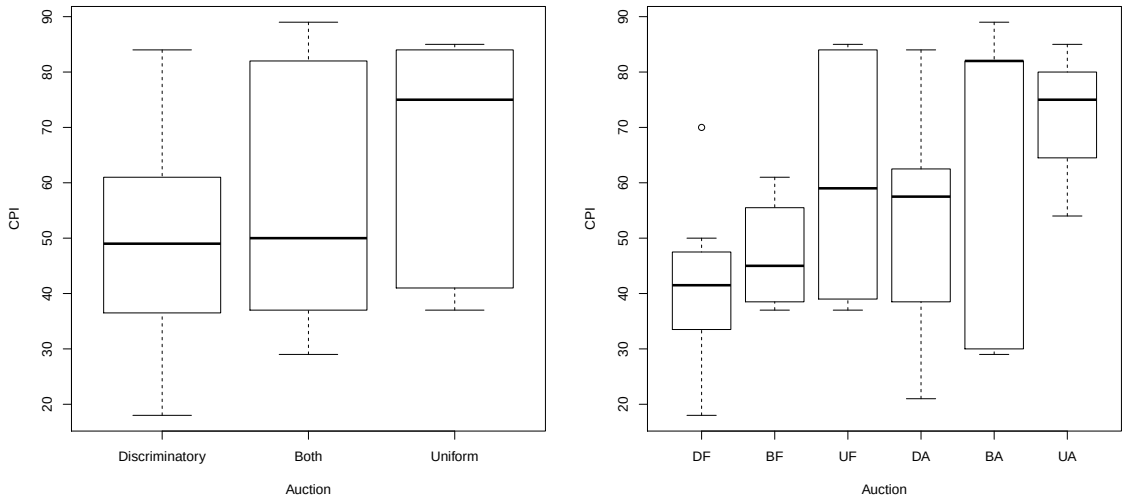

Figure 1: Corruption Perceptions Index 2017 (CPI) against auction format. Higher CPI indicates *lower* perceived corruption. In the left panel, auctions are binned by pricing method. In the right panel, auctions are binned by pricing method and whether supply is fixed (F) or adjustable (A). Concordant with Proposition 1, corruption is lower (CPI is higher) in nations which implement uniform price auctions than in countries that implement discriminatory auctions. In line with this hypothesis but outside of our current scope, corruption is lower in nations which commit to supply beforehand than in those which allow supply to depend on submitted bids.

4 Application: Sovereign Debt

Following Proposition 2 we now apply the hierarchy of auditability in multi-unit auctions to compare mechanism selection against perceived corruption. We use the classification of sovereign debt sales given in Brenner et al. [2009], in which each nation is categorized as using a discriminatory auction, a uniform-price auction, or some combination of the two to sell treasury securities.[21] Nations that utilize a combination of the two auctions implement one format or the other on a deterministic basis, so we make the assumption that their formats are of intermediate auditability, between the deterministic and uniform-price auctions. We use the 2017 Corruption Perceptions Index as our measure of perceived corruption.

Our results are displayed graphically in Figure 1, and regression results are in Table 1. These results are in line with the hypothesis that nations that are perceived to be more corrupt implement more auditable auctions.[22] Auditability makes corruption more difficult, and thus ceteris

[21] Nations such as Spain that implement neither format are excluded from the analysis.

[22] A more thorough analysis, controlling for features such as GDP, legal structure, and market competitiveness, is performed by Brenner et al. [2009]. Their results are directionally the same but statistically insignificant. Because

	Dependent variable:	
	CPI2017	
	(1)	(2)
Auction, Both	5.639	7.060
	(7.770)	(7.493)
Auction, Uniform	14.194*	18.458**
	(7.770)	(7.740)
Quantity, Adjustable		12.791**
		(6.175)
Constant	49.917***	41.390***
	(4.058)	(5.669)
Observations	42	42
R^2	0.080	0.173
Adjusted R^2	0.033	0.108
Residual Std. Error	19.878 (df = 39)	19.090 (df = 38)
F Statistic	1.697 (df = 2; 39)	2.657* (df = 3; 38)
Note:		*p<0.1; **p<0.05; ***p<0.01

Table 1: Results of regressing Corruption Perceptions Index on dummies for auction format and supply commitment.

paribus auditability makes mechanism participation more likely. In this sense, auditability and non-corruption are substitutes. Relatively noncorrupt nations have some latitude in mechanism selection, while relatively corrupt nations face strong incentives to auditability mechanisms.[23] Importantly, a mechanism need not be auditable to ensure participation: participants may have faith that the auction is run as claimed. This emphasizes the value of relative auditability. Fixing bidder credulity and public information release, there is a set of auctions which is sufficiently auditable for participation, and a set which is not.

While measured corruption may legitimize fears of corrupt dealing on the part of the auctioneer, to the extent corruption is a societal phenomenon it will also affect the auctioneer's decision of which auction to implement. Releasing public information about private bids simplifies collusion between bidders by making it easier to detect deviation from collusive agreements. Holding fixed

our goal is to expose the potential role of auditability, our results are complementary.

[23]This analysis ignores other reasons, such as revenue, that nations might implement one mechanism or another. Pycia and Woodward [2019] suggest that discriminatory auctions may revenue-dominate uniform price auctions, implying that a nation concerned with credibility *or* raising funds should implement a discriminatory auction. Other results (c.f. Ausubel et al. [2014], among many others) find ambiguous revenue comparisons that vary with context, implying that credibility may point in the opposite direction of other concerns.

the credulity of bidders, a less auditable mechanism requires more public information release to encourage bidder participation. For this reason we should expect auctioneers to implement more auditable mechanisms when collusion is more likely, as when a society has more corruption or there is greater market power. The former effect is identical to our hypothesis about more corrupt nations implementing more auditable mechanisms, and the latter is documented in Brenner et al. [2009].[24]

5 Conclusion

In this paper we have introduced the concept of auditability. Public information following an auction audits the auction's outcome if each bidder's outcome is consistent with the assumption that the seller is maximizing profits, conditional on consistency with bidder-private and public information. We obtain a natural hierarchy of auditability, and show that a price menu can be audited with no information while standard multi-unit auction formats require at least some information to audit their outcomes. Among common multi-unit auction formats, the discriminatory auction is more auditable than uniform price auction formats.

We extend this notion to include constrained auditability. When the seller can commit to a particular quantity (as with single-unit auctions of idiosyncratic items) the discriminatory auction is also auditable with no information. If the seller must announce the maximum market-clearing price, the discriminatory auction is equally as auditable as a last accepted bid uniform price auction, both of which are strictly more auditable than a first rejected bid uniform price auction. Under constraints on information and quantity commitment that roughly match practical implementations, there is a strict ordering: discriminatory auctions are more auditable than uniform price auctions, and last accepted bid auctions are more auditable than first rejected bid auctions.

In a brief empirical exercise we verify that corruption is positively correlated with the implementation of auditable auction formats, with regard to the sale of treasury securities. This is consistent with a treasury needing to assuage bidder misgivings about commitment while limiting information release.

[24] Brenner et al. [2009] associate competitiveness with a measure of economic freedom, or ease of doing business. Ease of doing business depends on the regulatory environment, which could affect the relative auditability of a mechanism by affecting the underlying informational environment. We ignore that effect here.

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.

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. Uniform price auctions with a last accepted bid pricing rule. 2018.

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.

Frank McSherry and Kunal Talwar. Mechanism design via differential privacy. In *Foundations of Computer Science, 2007. FOCS'07. 48th Annual IEEE Symposium on*, pages 94–103. IEEE, 2007.

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. Pay-as-bid: Selling divisible goods. 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.