Optimal Clock Auctions*

Simon Loertscher[†] Leslie M. Marx[‡]

May 13, 2018

Abstract

Clock auctions are weakly group strategy-proof, make bidding truthfully an obviously dominant strategy, and preserve trading agents' privacy. They have proved useful in practice but challenging to implement in a prior-free, asymptotically optimal way. We characterize the Bayesian optimal clock auction (BOCA) and develop a prior-free clock auction that maintains the structure of the BOCA and is asymptotically optimal. To do this, we exploit a relationship between hazard rates and the spacings between order statistics. Extensions permit price discrimination among heterogeneous groups, minimum revenue thresholds, and quantity caps.

Keywords: asymptotic optimality, Bayesian optimality absent estimation error, estimating virtual types, spacings

JEL Classification: C72, D44, L13

*We thank Salvador Barberà, Matt Jackson, Shengwu Li, Claudio Mezzetti, Paul Milgrom, Ellen Muir, Andras Niedermayer, Stefan Reichelstein, Larry Samuelson, Mark Satterthwaite, Ilya Segal, Thomas Tröger, Rakesh Vohra, Steve Williams, Robert Wolpert, and seminar participants at the Australian National University, Northwestern University, University of Florida, University of Melbourne, University of Rochester, University of Sydney, University of Technology in Sydney, University of Texas, the 2015 Decentralization Conference at Northwestern, the 2015 MaCCI Summer Institute, the 2015 EARIE Conference in Munich, the 2015 IO Theory Conference at Columbia, the 2016 Workshop on Auctions and Market Design at the University of Vienna, and the 2017 Australasian Economic Theory Workshop in Auckland for helpful comments. Support by the Faculty of Business and Economics at the University of Melbourne via an Eminent Research Scholar grant is gratefully acknowledged.

[†]Department of Economics, Level 4, FBE Building, 111 Barry Street, University of Melbourne, Victoria 3010, Australia. Email: simonl@unimelb.edu.au.

[‡]The Fuqua School of Business, Duke University, 100 Fuqua Drive, Durham, NC 27708, USA: Email: marx@duke.edu.

1 Introduction

Clock auctions have a number of properties that make them attractive for practical purposes. They are weakly group strategy-proof, preserve the privacy of trading agents, endow single-unit traders with obviously dominant strategies, and limit the information that agents and the designer must acquire prior to the auction.[1] Privacy preservation protects traders from hold-up by the designer and the designer from the (often political) risk of regret.[2] By endowing agents with dominant strategies, clock auctions exhibit equilibrium behavior that does not depend on common knowledge or higher-order beliefs. Therefore, they satisfy the robustness requirements emphasized by, for example, Bergemann and Morris (2012).[3]

However, clock auctions with optimally chosen reserve prices and stopping rules depend on the fine details of the environment and so are subject to what has become known as the Wilson critique (Wilson, 1987). Although there is a large economics and computer science literature on asymptotically optimal, prior-free mechanisms, to date none of these mechanisms is implementable as a clock auction. This creates a tension between prior-free, asymptotically optimal mechanisms that are not clock implementable and prior-free clock auctions that are not asymptotically optimal, seemingly leaving designers with the tough choice between one or the other.[4]

The tension is easily understood. In general, whether it is optimal for an agent to trade depends not only on his virtual type, but also on his ranking relative to agents on his side of the market, which is determined using the bids of the other players on his side of the market. This is the case in two-sided environments when the designer acts as an intermediary and also in one-sided auctions when the designer has a capacity constraint or otherwise increasing marginal costs. Even if the distribution that is used to gauge an agent's own virtual type does not depend on that agent's report, using his report to determine other agents' virtual types and their rankings relative to his may indirectly introduce a means to manipulate the

[1]The notion of *obviously dominant* strategies is defined by Li (2017). Li also shows that clock auctions have an equilibrium in obviously dominant strategies and that this implies weak group strategy-proofness. The point about limited information acquisition by traders is due to Milgrom and Segal (2015).

[2]Lucking-Reiley (2000) discusses hold-up by dealers of collectable stamps using second-price auctions and how truthful bidding was no longer a dominant strategy. Ex post regret was an issue following New Zealand's 1990 radio spectrum auction, which used a direct mechanism that revealed to the public the amount of money left on the table (McMillan, 1994; Milgrom, 2004).

[3]To the extent that clock auctions raise concerns, those relate to combinatorial clock auctions (see Levin and Skrzypacz, 2016) and so are not relevant here.

[4]Revenue extraction is often an important and sometimes the only design objective. For the pivotal role revenue considerations played in the U.S. Congress' decision to legislate the FCC to use auctions to allocate radio spectrum licenses, see for example Loertscher et al. (2015). Similarly, the question how much revenue the U.S. government should extract from the "incentive auction" was the subject of at times controversial debates. Even in economic theory, revenue plays an important role: The impossibility results of Vickrey (1961) and Myerson and Satterthwaite (1983) and their generalizations arise because the designer faces the constraint that revenue must not be negative.

mechanism.[5]

In this paper, we show how to reconcile prior-free clock auctions and asymptotic optimality. We exploit the insight that Myerson's *theoretical* construct of virtual types is tightly connected to the order statistics for nontrading agents' types and the spacings (distances) between them. In clock auctions, the spacings between nontrading agents' types are *observable* and, as we show, can be used to estimate the virtual types associated with the marginal active buyer and seller if agents on each side of the market draw their types from identical distributions. Because the estimates use only the reports of agents who do not trade, the privacy of the agents who trade is preserved. What is more, privacy preservation for trading agents guarantees incentive compatibility because, for example, buyers with higher values cannot influence the estimates and hence the ranks of lower valuing buyers, conditional on being active. Under the regularity assumption that virtual types are monotone, no knowledge of the inframarginal virtual types is required to determine the Bayesian optimal allocation.[6]

As we show, our prior-free clock auction is optimal in the sense that in the absence of estimation error it replicates the Bayesian optimal clock auction, which we characterize, and that the effects of estimation error vanish in the limit, so that the mechanism is asymptotically optimal. Although existing prior-free mechanisms are optimal absent estimation error for special cases such as constant marginal costs, our paper is the first to develop a mechanism that is optimal absent estimation error for a general setting and to formalize and apply this optimality criterion. Furthermore, our prior-free mechanism has a clock implementation whose structure is essentially uniquely pinned down by the requirement that it be sequentially consistent in that, similar to Akbarpour and Li's (2017) notion of credible mechanisms, there is no commitment problem for the auctioneer in the dynamic implementation. Estimation details can be further pinned down using criteria such as minimizing the mean square error of the estimator.

By extending our setup to have identifiable groups of buyers and groups of sellers, where agents are homogenous within groups but heterogeneous across groups, we can allow for price discrimination across groups, revenue thresholds, group-specific caps, and group-specific

[5]Goldberg et al. (2001) and Baliga and Vohra (2003) circumvent the problem by splitting the market into two sub-markets and using the estimates from one sub-market to determine the mechanism to be applied in the other. For the special case of constant marginal cost, Segal (2003) observes that one can use the empirical distribution based on other agents' reports without having to rely on estimates of the "empirical density" to determine whether a given buyer should trade.

[6]The inframarginal agents when there are k traders are the $k-1$ most efficient traders. Because clock auctions allocate the quantity traded to the most efficient traders, Bayesian optimality with finitely many agents cannot be implemented via a clock auction when virtual types are not monotone because optimality in that setting requires ironing (and hence an inefficient allocation with positive probability). Nevertheless, as we show, prior-free clock implementation of the Bayesian optimal mechanism is possible asymptotically even when virtual types are not monotone, provided price posting is Bayesian optimal in the large and there is a unique local maximum under price posting.

favoritism.

This paper contributes to the literature on clock auctions. Beginning with Milgrom and Weber (1982), with subsequent contributions by McAfee (1992), Kagel (1995), Lopomo (1998, 2000), Ausubel (2004, 2006), Milgrom and Segal (2015), and Li (2017), this literature has identified advantages of dynamic implementation over direct mechanisms in a variety of setups.[7] In particular, our paper builds on the properties of clock auctions identified by Milgrom and Segal (2015) and design features first introduced by McAfee (1992).

Motivated by Wilson (1987) and the literature on robust mechanism design in the tradition of Bergemann and Morris (2005, 2009, 2012),[8] we develop prior-free clock auctions that are Bayesian optimal absent estimation error and asymptotically optimal and thus lie at the intersection of robust and Bayesian mechanism design in the tradition of Myerson (1981), both for one-sided setups and for two-sided exchanges such as Myerson and Satterthwaite (1983), Gresik and Satterthwaite (1989), and Williams (1999). For a two-sided setting with multi-unit traders, Loertscher and Mezzetti (2016) develop a prior-free incentive-compatible clock auction in which the role for estimation is to gauge market demand and supply for the purpose of allocating efficiently without running a deficit.[9]

In the literature on asymptotically optimal, prior-free mechanisms, the two most important precursors to the current paper are Segal (2003) and Baliga and Vohra (2003).[10] Segal derives an asymptotically Bayesian optimal mechanism for one-sided setups when the designer is uncertain about the distribution of types but has a prior belief regarding the distribution. Baliga and Vohra (2003) construct dominant strategy prior-free mechanisms for one-sided and two-sided setups and show that in the limit with infinitely many traders, these mechanisms generate the same revenue as the Bayesian optimal mechanisms. Baliga and Vohra divide agents on each side of the market randomly into two groups and use reports from one group to estimate the virtual type functions for the other group.

Dominant strategy prior-free mechanisms have also received attention in the computer

[7]Although both the English auction and the second-price auction make bidding truthfully a dominant strategy, in laboratory settings subjects are consistently more likely to play their dominant strategy in English auctions than in second-price auctions (Kagel, 1995), suggesting that the open format of the English auction facilitates discovery of the dominant strategies.

[8]Hagerty and Rogerson (1987) provide an additional, related motivation for detail-free mechanisms: Environments are often subject to shocks while institutions that govern trade are longer-term in nature and must therefore be robust with respect to the details of changing environments.

[9]Because efficiency is a distribution-free concept, statistical properties in the setting of Loertscher and Mezzetti (2016) only matter for convergence, which allows them to depart from the independence assumption.

[10]There is also a vast literature on estimation in auctions using kernel density estimators; see Athey and Haile (2007) and Guerre et al. (2000) and the references therein. This literature is related because objects of interest here and there are the distributions from which bidders draw their types. Our k-th nearest neighbor estimator is a kernel density estimator based on the uniform kernel. In clock auctions, the estimates of interest are at the bound of the observed data, i.e., at a single point. For estimating densities at a single point, there appears to be no advantage of kernel estimates over nearest neighbor estimates (Silverman, 1986, pp. 20 and 97).

science literature. That literature analyzes mechanisms that use reports from a sample of agents to infer the distribution of types for other agents, referred to as random-sampling mechanisms.[11] Whereas the analysis of this type of mechanism in Baliga and Vohra (2003) focuses on profit maximization for the designer, the literature on Algorithmic Game Theory focuses on whether the mechanisms have good worst-case performance relative to benchmarks based on prior-free mechanisms that approximate Bayesian optimality but are not incentive compatible.[12] For example, Goldberg et al. (2001) and Dhangwatnotai et al. (2015) focus on the worst-case performance one-sided auctions for a good with unlimited supply, while Deshmukh et al. (2002) consider two-sided mechanisms.[13] None of these random-sampling mechanisms can be implemented as a clock auction.

This paper also relates to the large literature on micro-foundations for Walrasian equilibrium, whose modern guise goes back to Arrow (1959), Vickrey (1961), and Hurwicz (1973). How can a market maker infer the data necessary to clear the market without violating agents' incentive and participation constraints at no cost to himself? The short answer is that he cannot. However, one way of interpreting the results in McAfee (1992), Rustichini et al. (1994), Cripps and Swinkels (2006), and Satterthwaite et al. (2015) is that there are practical mechanisms that approximate full efficiency quickly as the economy grows. We show that the market maker's objective can be maximized, asymptotically, even if the objective is to maximize revenue or a convex combination of revenue and social surplus, without any prior knowledge or assumptions about distributions beyond mild regularity conditions and independence.

The remainder of this paper is structured as follows. Section 2 describes the setup and the relevant notions of optimality, including Bayesian optimality and prior-free optimality. Section 3 shows that in a two-sided setup, the Bayesian optimal mechanism is not clock implementable and derives the Bayesian optimal clock auction. Section 4 shows by construction that a prior-free clock auction exists that is prior-free optimal. In addition, we show that the structure of the prior-free optimal clock auction is pinned down by a notion of sequential consistency, and we provide criteria for determining the details of the required estimators. Section 5 contains extensions, and Section 6 concludes.

2 Setup

In this section, we define Bayesian, Bayesian optimal, and prior-free mechanisms, and we introduce a notion of optimality for prior-free mechanisms. We then specialize to the trade

[11] These mechanisms are referred to as "random sampling mechanisms" in, e.g., Goldberg et al. (2001) and Goldberg et al. (2006), but as "adaptive mechanisms" in Baliga and Vohra (2003).

[12] Devanur et al. (2015) provide a formal definition of "approximate" in this sense.

[13] Dütting et al. (2017) analyze two-sided mechanisms that can be implemented as clock auctions, but do not consider estimation.

settings that are the focus of this paper. For these settings, we describe the Bayesian optimal mechanisms along with the associated regularity requirements.

2.1 Bayesian and Bayesian optimal mechanisms

With little overhead cost in terms of notation, the following concepts can be introduced and discussed with a fair degree of generality. Let Θ_i be agent i's type space and $\Theta = \times_{i\in\mathbb{I}}\Theta_i$, where $\mathbb{I}$ is the set of agents, and let μ be a probability measure over Θ. A *direct mechanism* collects reports $\boldsymbol{\theta}$ and as a function of these reports determines individual quantities and transfers $\langle q_i(\boldsymbol{\theta},\mu), t_i(\boldsymbol{\theta},\mu)\rangle$ (we focus on deterministic mechanisms, which is without loss of generality under regularity assumptions imposed later). A direct mechanism is *incentive compatible* of some kind (e.g., Bayesian, dominant strategy or ex post) if for all i and all $\theta_i \in \Theta_i$, reporting truthfully constitutes an equilibrium of this kind. It is *individually rational* if the participation constraint of that kind is satisfied for all agents.

We say that a direct mechanism is *Bayesian* if there exist two priors μ and μ' with $\mu \neq \mu'$ and some $\boldsymbol{\theta} \in \Theta$ such that

$$\langle q_i(\boldsymbol{\theta},\mu), t_i(\boldsymbol{\theta},\mu)\rangle \neq \langle q_i(\boldsymbol{\theta},\mu'), t_i(\boldsymbol{\theta},\mu')\rangle.$$

Observe that whether a mechanism is Bayesian is independent of the nature of incentive compatibility.

We say that an incentive compatible and individually rational Bayesian mechanism $\langle q_i^*(\boldsymbol{\theta},\mu), t_i^*(\boldsymbol{\theta},\mu)\rangle$ is *Bayesian optimal* (or *optimal*) *with respect to prior* μ if for all $\boldsymbol{\theta} \in \Theta$,

$$\langle q_i^*(\boldsymbol{\theta},\mu), t_i^*(\boldsymbol{\theta},\mu)\rangle \in \arg\max_{\langle q_i,t_i\rangle} E_{\tilde{\boldsymbol{\theta}}|\mu}[W(\tilde{\boldsymbol{\theta}}, \langle q_i, t_i\rangle)], \tag{1}$$

where the maximization is subject to incentive compatibility and individual rationality constraints (and possibly additional constraints) and $E_{\tilde{\boldsymbol{\theta}}|\mu}[W(\tilde{\boldsymbol{\theta}}, \langle q_i, t_i\rangle)]$ denotes the expectation of the designer's objective W, which depends on the realized type profile and the mechanism, with the expectation being taken with respect to the random variable $\tilde{\boldsymbol{\theta}}$ with measure μ. When we refer to the Bayesian optimal mechanism without specifying the beliefs, we mean the Bayesian optimal mechanism with respect to correct beliefs. Bayesian optimality is particularly useful as a benchmark when it is known what this optimum is. This is the case in our setup with single-dimensional types, which we introduce below.

Because ex post efficiency is a distribution-free concept, we adhere to the convention of distinguishing between Bayesian optimal and efficient mechanisms. For example, Myerson (1981) derives optimal auctions whereas the Vickrey auction with a reserve equal to the seller's cost is an efficient auction.

2.2 Prior-free and prior-free optimal mechanisms

We say that a mechanism $\langle q_i, t_i \rangle$ is *prior free* if it is not Bayesian. That is, if for all priors μ and μ' and all $\boldsymbol{\theta} \in \Theta$,

$$\langle q_i(\boldsymbol{\theta}, \mu), t_i(\boldsymbol{\theta}, \mu) \rangle = \langle q_i(\boldsymbol{\theta}, \mu'), t_i(\boldsymbol{\theta}, \mu') \rangle,$$

in which case we can simply write the mechanism as $\langle q_i(\boldsymbol{\theta}), t_i(\boldsymbol{\theta}) \rangle$. Of course, a prior-free mechanism is detail free insofar as it does not depend on such things as the distributions from which agents draw their types.[14]

We say that a prior-free (and incentive-compatible and individually rational) mechanism $\langle q_i(\boldsymbol{\theta}), t_i(\boldsymbol{\theta}) \rangle$ is *Bayesian optimal absent estimation error (BOAEE)* if there exists a belief $\mu_{\boldsymbol{\theta}}$ that is "estimated," as indicated by its dependence on the reports $\boldsymbol{\theta}$**,** such that for all $\boldsymbol{\theta} \in \Theta$,

$$\langle q_i(\boldsymbol{\theta}), t_i(\boldsymbol{\theta}) \rangle = \langle q_i^*(\boldsymbol{\theta}, \mu_{\boldsymbol{\theta}}), t_i^*(\boldsymbol{\theta}, \mu_{\boldsymbol{\theta}}) \rangle.$$

That is, a prior-free mechanism is BOAEE if for all possible reports $\boldsymbol{\theta}$, it coincides with the Bayesian optimal mechanism for a designer with prior $\mu_{\boldsymbol{\theta}}$.[15]

Because BOAEE does not require the belief estimates to converge to the truth, the BOAEE property is most compelling when it is satisfied by a prior-free mechanism whose estimates do converge and whose estimator satisfies criteria such as minimizing mean square error. Loosely, we say that a prior-free mechanism is asymptotically optimal if, as the numbers of agents go to infinity (through the replication of an initial set of buyers and sellers), the expected value of the ratio of the objective under the mechanism to the objective under the Bayesian optimal mechanism with correct beliefs converges in probability to one.

Formally, in the tradition of Gresik and Satterthwaite (1989), we derive asymptotic results assuming independent private values by considering η-fold replicas of the economy. An η-fold replica has set of agents equal to the union of η instances of set $\mathbb{I}$ with the corresponding type space Θ for each replica, and beliefs defined over a single replica of the type space. We say that an incentive compatible, individually rational, prior-free mechanism $\langle q_i(\boldsymbol{\theta}; \eta), t_i(\boldsymbol{\theta}; \eta) \rangle$ defined for an η-fold replica is *asymptotically optimal* if

$$\text{plim}_{\eta \to \infty} E_{\boldsymbol{\theta}|\mu^*} \left[\frac{W(\boldsymbol{\theta}, \langle q_i(\boldsymbol{\theta}; \eta), t_i(\boldsymbol{\theta}; \eta) \rangle)}{W(\boldsymbol{\theta}, \langle q_i^*(\boldsymbol{\theta}, \mu^*; \eta), t_i^*(\boldsymbol{\theta}, \mu^*; \eta) \rangle)} \right] = 1,$$

where $\langle q_i^*(\boldsymbol{\theta}, \mu^*; \eta), t_i^*(\boldsymbol{\theta}, \mu^*; \eta) \rangle$ is the Bayesian optimal mechanism for the η-fold replica with

[14]The converse, interestingly, is not true. The first-price auction is generally considered detail free. For example, its rules make no references to distributions. Yet, it is not prior free because the Bayes-Nash equilibrium allocation varies with distributions. In particular, it is ex post efficient if and only if distributions are symmetric.

[15]Defined directly, $\langle q_i(\boldsymbol{\theta}), t_i(\boldsymbol{\theta}) \rangle \in \arg\max_{\langle q_i, t_i \rangle} E_{\tilde{\boldsymbol{\theta}}|\mu_{\boldsymbol{\theta}}}[W(\tilde{\boldsymbol{\theta}}, \langle q_i, t_i \rangle)]$, subject to incentive compatibility and individual rationality constraints.

respect to the correct beliefs and "plim" stands for convergence in probability.

We combine BOAEE and asymptotic optimality to define our notion of optimality for prior-free mechanisms. We say that a prior-free mechanism is *prior-free optimal* if it is BOAEE and asymptotically optimal.

2.3 Our trade settings

While the above concepts apply generally, our focus in this paper is on a narrower set of trading problems. In particular, we study settings in which the demand side is characterized by a vector of marginal valuations $\mathbf{v}$ of dimension n and the supply side by a vector of marginal costs $\mathbf{c}$ of dimension m. Letting $v_{(k)}$ and $c_{[k]}$ denote, respectively, the k-th highest and k-th lowest elements of $\mathbf{v}$ and $\mathbf{c}$, the efficient quantity traded is the largest integer k such that $v_{(k)} \geq c_{[k]}$, which is well defined using the conventions that $v_{(0)} \equiv \infty$, $v_{(n+1)} \equiv -\infty$, $c_{[0]} \equiv -\infty$, and $c_{[m+1]} \equiv \infty$. All trade occurs via a monopoly market maker, who is a risk-neutral designer without private information.

To account for private information, we assume that agents on at least one side of the market are privately informed about their types. Thus, we focus on setups with one-sided private information pertaining to buyers, one-sided private information pertaining to sellers, and two-sided private information.

We assume that all buyers draw their types from the same distribution and that all sellers draw their types from the same distribution, although buyer and seller distributions can differ. In the online appendix we extend our results allow for heterogeneity of buyers and heterogeneity among sellers.[16]

When considering prior-free mechanisms, we assume that each agent with private information is privately informed about his type, but the types and distributions from which they are drawn are unknown to the mechanism designer and to the agents. Keeping fixed the mechanism, equilibrium behavior would not be affected if the agents knew the distributions because the mechanism endows them with dominant strategies; however, depending on assumptions about the informational structure, alternative mechanisms, such as those developed by Crémer and McLean (1985, 1988) could be optimal.[17] We assume that the designer only knows that buyers and sellers draw their types independently from the same distributions and that the Bayesian design problem satisfies certain other conditions that we

[16] Specifically, in the extension we assume that buyers can be placed into groups, where all buyers within a group draw their values from the same distribution, and that sellers can be placed into groups, where all sellers within a group draw their costs from the same distribution. But we allow heterogeneity across groups. This extension allows for price discrimination across groups as well as the straightforward implementation of revenue thresholds, group-specific caps, and group-specific favoritism.

[17] However, Crémer-McLean mechanisms are not prior-free mechanisms as defined here. While the allocation rule of the full-surplus extracting mechanism is ex post efficient and thus independent of the prior, the ex post transfers vary with distributions.

spell out below. For clock implementation, we assume that the designer knows upper and lower bounds for the types (not necessarily tight ones), which allows the designer to start the clock auction at prices that guarantee that all agents are active irrespective of their types.

For prior-free mechanisms, dominant strategy incentive compatibility seems like the natural notion of incentive compatibility. Moreover, it is without loss of generality when there is no restriction on the set of admissible priors.[18] Individual rationality is most naturally required to be satisfied ex post.[19] These are therefore the notions we focus on going forward.

2.4 Benchmark Bayesian optimal mechanisms

We now describe the Bayesian optimal mechanisms for the informational setups that we consider as well as regularity assumptions. The Bayesian optimal mechanisms do not require knowledge of the virtual types of the inframarginal traders to determine the quantity traded or the payments in the dominant strategy implementation. For the prior-free setting, this implies that the virtual types of the inframarginal types need not be estimated. As will be seen shortly, this opens the scope for clock implementation.

Bayesian optimality with one-sided private information pertaining to buyers

For the setup with *one-sided private information pertaining to buyers*, $c_{[1]}, ..., c_{[m]}$ defines the commonly known marginal cost curve of the designer, and there are n buyers who have unit demands and draw their valuations independently from the continuously differentiable distribution function F with support $[\underline{v}, \overline{v}]$ and positive density f everywhere on the support.[20] In the above notation, the type space is $\Theta = [\underline{v}, \overline{v}]^n$, with correct beliefs defined with respect to F. A buyer's payoff is equal to his value minus the price he pays if he trades and zero otherwise. It is well known that under the assumptions that $\underline{v} \leq c_{[1]} < \overline{v}$ and that the virtual valuation function

$$\Phi(v) \equiv v - \frac{1 - F(v)}{f(v)}$$

[18] To see this, recall first from Bergemann and Morris (2005) that for private values environments like ours dominant strategy implementation is equivalent to ex post implementation. Therefore, Bayesian incentive compatibility is the only alternative notion of incentive compatibility. For the purpose of reaching a contradiction, assume then that a mechanism satisfies Bayesian incentive compatibility but fails to be dominant strategy incentive compatible. Because the mechanism is prior-free, it must be Bayesian incentive compatible for any admissible prior, including priors with mass one at the type profile(s) for which the mechanism fails to satisfy dominant strategy incentive compatibility. But for such priors, Bayesian incentive compatibility reduces to dominant strategy incentive compatibility, which is the desired contradiction.

[19] In Bayesian mechanism design settings with private values, there is an equivalence between the Bayesian and dominant strategy notions of incentive compatibility and of interim and ex post individual rationality; see, for example, Manelli and Vincent (2010) and the generalization by Gershkov et al. (2013).

[20] We assume continuously differentiable distributions instead of merely continuous distributions because our asymptotic results rely on the continuity of the inverses of the virtual type functions, which is guaranteed if the densities are continuous.

is increasing, the solution to the designer's profit maximization problem, which is subject to buyers' incentive compatibility and individual rationality constraints, has an allocation rule that trades the quantity given by the largest index q such that $\Phi(v_{(q)}) \geq c_{[q]}$.[21] All buyers with a value of at least $v_{(q)}$ trade and, in the dominant strategy implementation, pay the price $p^B = \max\{v_{(q+1)}, \Phi^{-1}(c_{[q]})\}$. This is a standard sales auction with a reserve that depends on the quantity traded.

Bayesian optimality with one-sided private information pertaining to sellers

Analogously, for *one-sided private information pertaining to sellers*, we assume that sellers have unit capacities and draw their privately known costs independently from a continuously differentiable distribution G with support $[\underline{c}, \overline{c}]$ and positive density g on the support. A seller's payoff is equal to the payment she receives minus her cost if she trades and zero otherwise. If the designer's marginal values $v_{(1)}, v_{(2)}, .., v_{(n)}$ are commonly known, $\overline{c} \geq v_{(1)} > \underline{c}$, and the virtual cost function

$$\Gamma(c) \equiv c + \frac{G(c)}{g(c)}$$

is increasing, then this is a standard procurement auction in which the optimal quantity traded is the largest index q such that $v_{(q)} \geq \Gamma(c_{[q]})$.[22] All sellers with costs not larger than $c_{[q]}$ trade and, in the dominant strategy implementation, are paid $p^S = \min\{c_{[q+1]}, \Gamma^{-1}(v_{(q)})\}$.

Bayesian optimality with two-sided private information

For *two-sided private information*, we let the set of (privately informed) agents be $\mathbb{I} = \mathbb{N} \cup \mathbb{M}$, where $\mathbb{N}$ is the set of buyers with unit demands, whose cardinality is n, and $\mathbb{M}$ with cardinality m is the set of sellers with unit capacities. As above, buyers and sellers have quasi-linear payoffs and outside options of value zero. Our problem is most interesting when, under the optimal Bayesian mechanism, full trade is sometimes but not always optimal, with full trade meaning that the quantity traded is $\min\{n, m\}$. A simple condition that guarantees this for the setting with two-sided private information is

$$\overline{c} \geq \overline{v} > \underline{c} \geq \underline{v}. \tag{2}$$

[21] After accounting for incentive compatibility and individual rationality constraints in a direct mechanism, the designer's problem is to choose a feasible allocation rule to maximize $E_{\mathbf{v}|F,\ldots,F}[\sum_{i=1}^{n} \Phi(v_i) q_i(\mathbf{v}, \mathbf{c}) - \sum_{j=1}^{m} c_j q_j(\mathbf{v}, \mathbf{c})]$, where q_i is the probability that buyer i receives a unit and q_j is probability that the j-th unit is produced, with feasibility meaning that $\sum_{i=1}^{n} q_i(\mathbf{v}, \mathbf{c}) \leq \sum_{j=1}^{m} q_j(\mathbf{v}, \mathbf{c})$.

[22] To see this, apply standard arguments to conclude that, after accounting for incentive compatibility and individual rationality constraints, the designer's problem in a direct mechanism is to choose a feasible allocation rule to maximize $E_{\mathbf{c}|G,\ldots,G}[\sum_{i=1}^{n} v_i q_i(\mathbf{v}, \mathbf{c}) - \sum_{j=1}^{m} \Gamma(c_j) q_j(\mathbf{v}, \mathbf{c})]$.

We refer to condition (2) and its one-sided analogues, $\underline{v} \leq c_{[1]} < \overline{v}$ and $\overline{c} \geq v_{(1)} > \underline{c}$, as *no-full trade* conditions, and throughout the paper we assume that the relevant no-full trade condition holds. Under this condition, assuming that Φ and Γ are increasing functions, the Bayesian optimal mechanism in the two-sided setting is characterized by the allocation rule that given $(\mathbf{v}, \mathbf{c})$ trades the quantity q that is the largest index such that $\Phi(v_{(q)}) \geq \Gamma(c_{[q]})$.[23] As in the one-sided setups, the number q is unique almost surely because ties among agents' types are a probability zero event. Buyers with values no less than $v_{(q)}$ trade and sellers with costs not larger than $c_{[q]}$ trade. In the dominant strategy implementation, trading buyers pay $p^B = \max\{v_{(q+1)}, \Phi^{-1}(\Gamma(c_{[q]}))\}$ and trading sellers are paid $p^S = \min\{c_{[q+1]}, \Gamma^{-1}(\Phi(v_{(q)}))\}$.

Regularity assumptions

Setups with monotonically increasing virtual type functions correspond to what Myerson (1981) refers to as the *regular case*. Thus, we say that the *regularity condition* is satisfied if the virtual type functions are monotonically increasing. Monotonicity of the virtual type functions ensures that point-by-point maximization permits incentive compatibility because it implies that more efficient types—buyers with higher values, sellers with lower costs—are more likely to trade. If, for a two-sided setting, Φ is not monotone in the neighborhood of some v' with $\Phi(v') > \underline{c}$ and/or Γ is not monotone in the neighborhood of some c' with $\Gamma(c') < \overline{v}$,[24] then the Bayesian optimal mechanism differs from the one described above. For finite numbers of buyers and sellers, there is ironing and random rationing with positive probability.

However, for asymptotic optimality, the regularity condition can be relaxed considerably. To illustrate, suppose a two-sided setup and confine attention, temporarily, to posted-price mechanisms. A necessary condition for the posted prices p^B and p^S to be profit maximizing in the large among all prices $(\hat{p}^B, \hat{p}^S) \in [\underline{v}, \overline{v}] \times [\underline{c}, \overline{c}]$ is that they satisfy

$$\Phi(p^B) = \Gamma(p^S) \quad \text{and} \quad n(1 - F(p^B)) = mG(p^S). \tag{3}$$

If the solution to (3) is unique and if price posting is the Bayesian optimal mechanism in the large (both of which are the case when the regularity condition is satisfied), then the mechanisms described above converge to the Bayesian optimal mechanisms even when they are not Bayesian optimal in the small.[25] For example, Figure 1 illustrates a case in which, in the large with equal numbers of buyers and sellers, the Bayesian optimum is

[23] In a direct incentive compatible and individual rational mechanism, standard arguments imply that the designer's problem reduces to choosing a feasible allocation rule to maximize $E_{\mathbf{v},\mathbf{c}|F,\ldots,F,G,\ldots,G}[\sum_{i=1}^{n} \Phi(v_i)q_i(\mathbf{v}, \mathbf{c}) - \sum_{j=1}^{m} \Gamma(c_j)q_j(\mathbf{v}, \mathbf{c})]$.

[24] When private information pertains only to buyers, respectively sellers, the conditions have to be replaced by $\Phi(v') > c_{[1]}$ and $\Gamma(c') < v_{(1)}$.

[25] When private information only pertains to buyers, respectively sellers, the optimal posted prices are obtained from (3) by replacing Γ and Φ by the identity function.

implemented with posted prices p^B and p^S such that share q^* of buyers and sellers trade, i.e., $p^B = F^{-1}(1-q^*)$ and $p^S = G^{-1}(q^*)$. However, in the small, because of ironing, random rationing occurs with positive probability and the mechanisms described above are not Bayesian optimal.

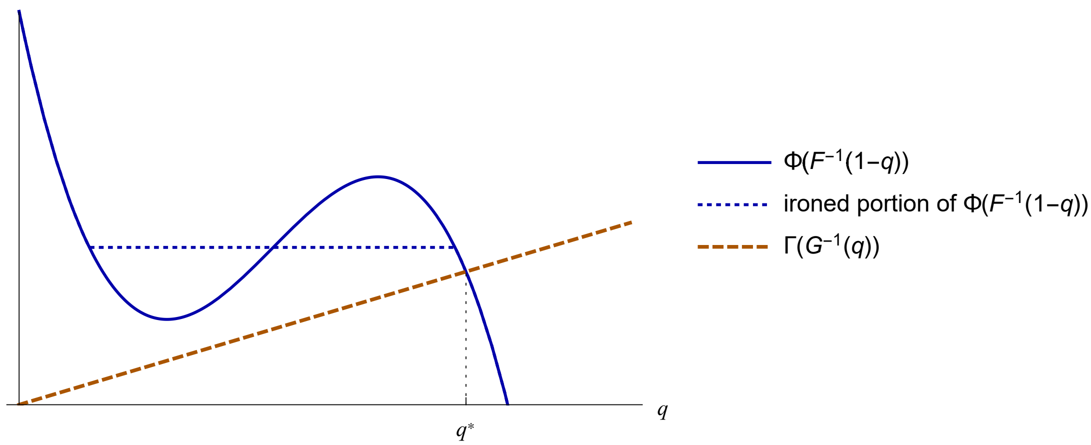

Figure 1: Illustration of a setup with unique posted prices that are Bayesian optimal in the large, but where ironing may be required in the small.

If the example of Figure 1 were adjusted so that the virtual cost function intersected the ironed portion of the virtual value function, then the Bayesian optimal mechanism in the large would no longer be a price posting mechanism because some share of agents would have to be rationed.

Our discussions and results below pertaining to Bayesian optimal mechanisms in the small and to BOAEE mechanisms assume that the regularity condition is satisfied. Results concerning asymptotic optimality hold under the weaker condition that there is a unique pair of prices satisfying (3) and that price posting is the Bayesian optimal mechanism in the large.

The tight connection between the Bayesian optimality of price posting in the large and the asymptotic optimality of our clock auction is not surprising; afterall, a clock auction generates prices that trading agents take as given and that are, in that sense, posted to them. The requirement that the prices p^B and p^S satisfying (3) be unique relates to the fact that, as we shall see, our clock auction stops at the first p^B and p^S such that (3) holds.

Generalized designer objective

The scope of the analysis can be further generalized by assuming that the designer wants to maximize a Ramsey objective, that is, a weighted sum of expected profit and social surplus,[26]

[26]Social surplus is defined to be the sum of trading buyers' values minus the sum of trading sellers' costs.

with weight $\alpha \in [0,1]$ on expected revenue, subject to incentive compatibility and individual rationality. The Bayesian optimal mechanism is then characterized by the same allocation rules as derived above for the case of profit maximization, except that one has to replace the virtual value and virtual cost functions by, respectively, the relevant *weighted virtual type function* $\Phi_\alpha(v)$ and $\Gamma_\alpha(c)$ defined as

$$\Phi_\alpha(v) \equiv \alpha\Phi(v) + (1-\alpha)v = v - \alpha\frac{1-F(v)}{f(v)} \quad \text{and} \quad \Gamma_\alpha(c) \equiv \alpha\Gamma(c) + (1-\alpha)c = c + \alpha\frac{G(c)}{g(c)}.$$

By construction, for $\alpha = 0$ the weighted virtual types correspond to true types, so that $\alpha = 0$ is equivalent to ex post efficiency, and for $\alpha = 1$, we have $\Phi(v) = \Phi_1(v)$ and $\Gamma(c) = \Gamma_1(c)$, so that $\alpha = 1$ corresponds to profit maximization.[27] The payments in the dominant strategy implementation are accordingly defined by replacing the virtual type functions Φ and Γ (and their inverses) by Φ_α and Γ_α (and their inverses) in the formulas above.

Beyond generality, allowing the designer to have a Ramsey objective also highlights the need for estimation because, as soon as $\alpha > 0$, the optimal mechanism depends on distributional details.

3 Bayesian optimal clock auctions

In this section we define clock auctions, discuss some of their key characteristics, and derive the Bayesian optimal clock auction.

3.1 Definition of a clock auction

In a clock auction, active buyers and sellers choose whether to exit as the buyer clock price increases and the seller clock price decreases, but agents who exit remain inactive thereafter. When the auction ends, active agents trade, with active buyers paying the buyer clock price and active sellers receiving the seller clock price.

To formally define a clock auction for our setup, we adapt the definition of a clock auction in Milgrom and Segal (2015) to accommodate (without requiring) a two-sided setting.[28] In

[27] As noted by Bulow and Roberts (1989), when $\alpha = 1$, the virtual values and virtual costs can be interpreted, respectively, as a buyer's marginal revenue and a seller's marginal cost, treating the (change in the) probability of trade as the (marginal change in) quantity. For $\alpha \in (0,1)$, the weighted virtual values and costs are convex combinations of the true and the virtual types, with weight α attached to the virtual types. If the social shadow cost of taxation, which is a measure of the distortion associated with raising revenue through taxes, is known to be some $\lambda \geq 0$, then α can be chosen to implement the socially optimal allocation by choosing $\alpha = \lambda/(1+\lambda)$ (see, e.g., Norman (2004) or Loertscher et al. (2015)).

[28] Milgrom and Segal (2015) define a (descending) clock auction for the one-sided setup. Their specification differs from ours in that it has individual-specific clocks and proceeds in discrete periods in which prices from a finite set are offered to the agents. Their clock auction is defined in terms of a price mapping from histories that are sequences of nested sets of active agents, where the price weakly decreases as agents exit

particular, to account for two sides, one needs to ensure that the numbers of active buyers and of active sellers are the same at the time the procedure ends.

In our setup, a *clock auction* is a rule for determining state transitions for a state space Ω, where the state $\boldsymbol{\omega}$ keeps track of: the number of active buyers and sellers, the exit prices of the nonactive buyers and sellers, the current buyer clock price, the current seller clock price, and whether the auction has ended. State transitions are governed by three functions: a buyer function $\phi : \Omega \to \mathcal{R}$, which is increasing in the buyer clock price, a seller function $\gamma : \Omega \to \mathcal{R}$, which is increasing in the seller clock price, and a target function $\tau : \Omega \to \mathcal{R}$, which satisfies $\tau(\boldsymbol{\omega}) \in [\phi(\boldsymbol{\omega}), \gamma(\boldsymbol{\omega})]$ whenever $\phi(\boldsymbol{\omega}) < \gamma(\boldsymbol{\omega})$. Because these three functions determine the state transitions, they also determine when the clock auction ends. As mentioned above, once the auction ends, the remaining active buyers buy at the buyer clock price, and the remaining active sellers sell at the seller clock price. Because a clock auction is defined by the functions ϕ, γ, and τ, we denote a clock auction by $\mathcal{C}_{\phi,\gamma,\tau}$. We provide the full definition of a clock auction for the two-sided setup (and the adaptation for a one-sided setup) in Appendix A.

In a clock auction, if there are unequal numbers of active buyers and sellers, then the clock price on the long side is advanced until exits on that side of the market equalize the number of active buyers and sellers. Once there are equal numbers of active buyers and sellers, the buyer and seller functions are evaluated at the current state. The auction ends if the state $\boldsymbol{\omega}$ is such that the value of the buyer function is greater than or equal to the value of the seller function, that is, $\phi(\boldsymbol{\omega}) \geq \gamma(\boldsymbol{\omega})$. If not, then the target function comes into play.

The value of the target function $\tau(\boldsymbol{\omega})$, which when $\phi(\boldsymbol{\omega}) < \gamma(\boldsymbol{\omega})$ is weakly between the values of the buyer and seller functions, becomes a target for the buyer and seller functions to achieve. Specifically, the buyer clock price is increased so as to increase the value of the buyer function towards $\tau(\boldsymbol{\omega})$, holding fixed the components of the state other than the buyer clock price, and the seller clock price is decreased so as to decrease the value of the seller function towards $\tau(\boldsymbol{\omega})$, holding fixed the components of the state other than the seller clock price.

If the target is reached on both the buyer side and the seller side with no exits (indicating that the value of the buyer function would have exceeded the value of the seller function following the next pair of exits), then the auction ends. Otherwise the auction continues, moving the clock price on the long side to equalize the numbers of active buyers and sellers, updating the state, and reevaluating the buyer and seller functions and the target function. By construction, the clock auction ends, either with trade or because all agents have exited.

As an example, McAfee's (1992) asymptotically efficient clock auction fits within our definition of a clock auction.[29] It corresponds to $\mathcal{C}_{\phi,\gamma,\tau}$ where, given a state $\boldsymbol{\omega}$ with equal

the active set.

[29]In McAfee's (1992) auction, the buyer clock price increases and seller clock price decreases, inducing exits by buyers and sellers. If there are unequal numbers of active buyers and sellers, then the clock price

numbers of buyers and sellers and clock prices p^B and p^S, the buyer and seller functions are essentially identity functions, with $\phi(\boldsymbol{\omega}) = p^B$ and $\gamma(\boldsymbol{\omega}) = p^S$, and the target function gives the midpoint between the two clock prices, $\tau(\boldsymbol{\omega}) = \frac{p^B+p^S}{2}$.

We assume that each buyer observes at least the buyer clock price and each seller observes at least the seller clock price. Agents' strategies are mappings from observed histories to exit decisions. The truthful strategy of a buyer is to exit if and only if the buyer clock price is greater than or equal to his value, and the truthful strategy of a seller is to exit if and only if the seller clock price is less than or equal to her cost. Playing these truthful strategies is dominant strategy incentive compatible.

3.2 Key properties of clock auctions

Clock auctions are well suited for practical implementation, and uniquely so on some dimensions.

In our environment with unit demands and unit supplies, Li (2017) shows that clock auctions, and only clock auctions, have *obviously dominant strategies*: the maximum payoff obtained by deviating from a dominant strategy at given price is never more than the minimum payoff obtained by sticking to the dominant strategy. Specifically, if a buyer exits before the buyer clock price reaches his value, his payoff is zero, and if a buyer remains active after the buyer clock price reaches his value, his payoff is bounded above by zero; however, under truthful bidding his payoff is bounded below by zero.

Given this property, it follows that clock auctions have dominant strategies (obviously) and are *weakly group strategy-proof*: for every profile of types, every subset of agents, and every deviant strategy profile for these agents, at least one agent in the subset has a weakly higher payoff from exiting when the clock price reaches the agent's type than from the deviant strategy profile.[30] In addition, a clock auction is *envy free* in the sense that in equilibrium no agent prefers the allocation and price of another agent to his own.[31]

on the long side is advanced in order to induce exit(s) and equalize the number of active buyers and seller. Following an exit that leaves equal numbers of active buyers and sellers, the auction ends if the last buyer exit occurred at a buyer clock price of p^B that is greater than or equal to the seller clock price p^S at which the last seller exited. Active agents then trade at their clock prices. (Thus, McAfee's stopping rule asks whether $p^B \geq p^S$ rather than comparing functions of p^B and p^S as in our auction.) If $p^B < p^S$, then the auction continues by setting a target, where in McAfee (1992), the target takes the form of a target price that both the buyer and seller clock prices move towards, defined as $\frac{p^B+p^S}{2}$. The buyer price is increased towards the target and the seller price is decreased towards the target. If the target is reached for both buyers and sellers with no exits, then the auction ends and active agents trade at the (common) clock price. Otherwise, the auction continues.

[30] On weak group strategy-proofness in a one-sided clock auction, see Li (2017) and Milgrom and Segal (2015). On the connection between individual and group strategy-proofness, see Barberà, Berga, and Moreno (2014). Dütting et al. (2017) show weak group strategy-proofness holds for a "lookback composition" of buyers and sellers that are ranked according to their types, which is a special case of the clock auctions considered here in that it has no target prices.

[31] This property is not unique to the clock auction implementation. The dominant strategy implementa-

These properties imply that clock auctions are robust with respect to the fine details of the environment and that, in the absence of transfers, collusion among a subset of agents cannot be strictly profitable for all of the colluding agents. Further, because endowing agents with dominant strategies and having agents recognize their dominant strategies are two distinct things in practice, the value of having obviously dominant strategies is a powerful argument for the use of clock auctions and for focusing on direct mechanisms that can be implemented via clock auctions. Indeed, this underlies the view expressed by Dasgupta and Maskin (2000) that the development of appropriate dynamic counterparts to Vickrey auctions is a leading topic for further research.

3.3 Implementation of Bayesian mechanisms by clock auction

As we have described, clock auctions possess a number of desirable properties. Thus, it is of interest when the Bayesian optimal mechanism can be implemented by a clock auction.

Proposition 1 *Assuming the regularity condition holds, in the setup with one-sided private information, the Bayesian optimal mechanism can be implemented by a clock auction; however, in the setup with two-sided private information, it cannot.*

Proposition 1 summarizes the implication of the results of Milgrom and Segal (2015) for our setting.[32] In one-sided settings, the Bayesian optimal mechanism can be defined based only on the information held by the designer and information gleaned from nontrading agents, and thus has a clock implementation.[33] In contrast, in two-sided settings, the Bayesian optimal mechanism relies on the private information of some trading agents—information that is not available in clock auctions because they preserve the privacy of trading agents. Thus, in two-sided settings, clock auctions do not always allow for the optimal quantity to be traded and so are with some loss of generality.

Given Proposition 1, in order to use the *Bayesian optimal clock auction (BOCA)* as a benchmark for a prior-free clock auction in a two-sided setting, we must first identify what the BOCA is in a two-sided setting. To do so, we temporarily stipulate that the designer knows F and G, and so knows the weighted virtual value function Φ_α and the weighted virtual cost function Γ_α, but is restricted to using a clock auction.

tions of the Bayesian optimal mechanisms derived above are also envy free.

[32]For settings like ours, Proposition 6 in Milgrom and Segal (2015) shows that clock implementation implies that agents are substitutes, while their Proposition 7 shows that clock implementation is possible when agents are substitutes. Our Proposition 1 then follows once one notices that, with private information pertaining to only one side of the market, the privately informed agents are substitutes to each other, whereas when private information pertains to both sides, buyers and sellers are complements—the problem is assignment representable, as described by Delacrétaz et al. (2018), and so the complementarity of buyers and sellers follows from Shapley (1962).

[33]For example, with private information only on the buyer side, the Bayesian optimal mechanism is implemented by the clock auction $\mathcal{C}_{\phi,\gamma,\tau}$ where, given a state $\boldsymbol{\omega}$ with buyer clock price p_B and n^A active buyers, $\phi(\boldsymbol{\omega}) \equiv \Phi_\alpha(p_B)$, $\gamma(\boldsymbol{\omega}) \equiv c_{[n^A+1]}$, and $\tau(\boldsymbol{\omega}) \equiv c_{[n^A]}$.

If the j-th highest valuing buyer and j-th lowest cost seller exit at a buyer clock price of $p^B = v_{(j)}$ and a seller clock price of $p^S = c_{[j]}$, then Bayesian optimality requires that the auction end if and only if $\Phi_\alpha(p^B) \geq \Gamma_\alpha(p^S)$. Thus, given a state $\boldsymbol{\omega}$ with current buyer and seller clock prices p^B and p^S, the BOCA must use a buyer function of $\phi(\boldsymbol{\omega}) = \Phi_\alpha(p^B)$ and a seller function of $\gamma(\boldsymbol{\omega}) = \Gamma_\alpha(p^S)$. The more subtle issue is how to set the target function.

If $\Phi_\alpha(v_{(j)}) < \Gamma_\alpha(c_{[j]})$ and $\Phi_\alpha(v_{(j-1)}) \geq \Gamma_\alpha(c_{[j-1]})$, then Bayesian optimality requires that the $j-1$ highest-valuing buyers and lowest-cost seller trade. However, for the two-sided setup, the privacy preservation inherent in a clock auction prevents this from always being achieved.[34] Given the constraint of privacy preservation for trading agents, the constrained optimum when it is observed that $\Phi_\alpha(v_{(j)}) < \Gamma_\alpha(c_{[j]})$ is to set target prices that maximize the probability that there are $j-1$ trades conditional on $\Phi_\alpha(v_{(j-1)}) \geq \Gamma_\alpha(c_{[j-1]})$. Thus, assuming that the regularity condition holds, in the two-sided setup, the BOCA is the clock auction $\mathcal{C}_{\phi,\gamma,\tau}$ such that for any state $\boldsymbol{\omega}$ with $j-1$ active buyers and sellers and current buyer and seller clock prices of p^B and p^S (equal to $v_{(j)}$ and $c_{[j]}$ under truthful bidding), the buyer and seller functions satisfy

$$\phi(\boldsymbol{\omega}) = \Phi_\alpha(p^B) \quad \text{and} \quad \gamma(\boldsymbol{\omega}) = \Gamma_\alpha(p^S), \tag{4}$$

and if $\Phi_\alpha(p^B) < \Gamma_\alpha(p^S)$, the target function maximizes the probability that the quantity traded is $j-1$ conditional on $\Phi_\alpha(v_{(j-1)}) \geq \Gamma_\alpha(c_{[j-1]})$, where the probability is taken with respect to random variables $v_{(j-1)}$ and $c_{[j-1]}$ given the observed values of $p^B = v_{(j)}$ and $p^S = c_{[j]}$, i.e.,

$$\tau(\boldsymbol{\omega}) \in \underset{\tau\in[\Phi_\alpha(p^B),\Gamma_\alpha(p^S)]}{\arg\max} \Pr \left(\Gamma_\alpha(c_{[j-1]})) \leq \tau \leq \Phi_\alpha(v_{(j-1)}) \mid v_{(j)} = p^B, c_{[j]} = p^S, \Gamma_\alpha(c_{[j-1]}) \leq \Phi_\alpha(v_{(j-1)})\right). \tag{5}$$

In other words, the BOCA uses the weighted virtual type functions as the buyer and seller functions and has a target function of the form described in (5).

As shown in the following proposition, under certain distributional assumptions we can further characterize the BOCA target function.

Proposition 2 *Assuming the regularity condition holds, if $f/(1-F)$ is increasing and concave and g/G is decreasing and concave, then in the two-sided setup, the BOCA target function is $\tau(\boldsymbol{\omega}) = \min\left\{\Gamma_\alpha(p^S), \max\left\{\Phi_\alpha(p^B), \delta^*\right\}\right\}$, where δ^* satisfies*

$$\frac{f(\Phi_\alpha^{-1}(\delta^*))}{1-F(\Phi_\alpha^{-1}(\delta^*))}\frac{1}{\Phi'_\alpha(\Phi_\alpha^{-1}(\delta^*))} = \frac{g(\Gamma_\alpha^{-1}(\delta^*))}{G(\Gamma_\alpha^{-1}(\delta^*))}\frac{1}{\Gamma'_\alpha(\Gamma_\alpha^{-1}(\delta^*))}. \tag{6}$$

Proof. See Appendix B.

[34] In the online appendix, we describe "quasi-clock auctions" that implement the Bayesian optimal mechanism in a two-sided setup without violating privacy preservation for any trading agents other than the marginal pair.

For example, under the conditions of Proposition 2, when $\alpha = 0$, then (6) reduces to $\frac{f(\delta^*)}{1-F(\delta^*)} = \frac{g(\delta^*)}{G(\delta^*)}$, which is to say that δ^* simply equalizes the hazard rates, which are assumed monotone. For any $\alpha \in [0,1]$, if F and G are uniform on $[0,1]$, then $\delta^* = 1/2$. This means that, in this example, whenever there are equal numbers of active buyers and sellers and the clock prices p^B and p^S are such that $\Phi_\alpha(p^B) < \Gamma_\alpha(p^S)$, then the BOCA simply increases the buyer clock price towards $\Phi_\alpha^{-1}(1/2)$ (or holds the buyer clock price fixed if it is already above this level) and decreases the seller clock price towards $\Gamma_\alpha^{-1}(1/2)$ (or holds the seller clock price fixed if it is already below this level). If both clocks obtain their respective target prices with no exits, then the auction ends.

4 Optimality for prior-free clock auctions

We now show the existence of a prior-free clock auction that is prior-free optimal. The requirement of BOAEE for a prior-free optimal clock auction identifies the mappings that need to be estimated, namely the buyer function, seller function, and target function. The requirement of asymptotic optimality places constraints on how that estimation is done. In addition, we show that the structure of the prior-free optimal clock auction is pinned down by a notion of sequential consistency, which we define below. Within that structure, we provide criteria for selecting the estimators to be used.

The exposition focuses on the case of two-sided private information. We discuss the adjustments required for the case of one-sided private information following Lemma 5.

Implications of BOAEE and asymptotic optimality

In order to have a prior-free clock auction that is prior-free optimal, it must be BOAEE for a clock auction and asymptotically optimal. In order to have BOAEE for a clock auction, we require prior-free estimates of the virtual type functions $\Phi_\alpha(\cdot)$ and $\Gamma_\alpha(\cdot)$ and the target virtual type that, for $j \in \{2, ..., \min\{m,n\}+1\}$, depend only on $\mathbf{v}_{(j)} \equiv (v_{(j)}, ..., v_{(n)})$ and $\mathbf{c}_{[j]} \equiv (c_{[j]}, ..., c_{[m]})$ because these are the only data that would be available in a clock auction. However, in our setup this requirement is not actually a restriction, as shown in the following lemma.[35]

Lemma 1 *For any dominant strategy incentive compatible, ex post individually rational, and envy-free direct mechanism, there exists an evaluation function family* $\{e_j\}_{j=0}^{\min\{n,m\}+1}$

[35] Goldberg et al. (2001) provide a similar result to Lemma 1 for the case of one-sided auctions for a good with unlimited supply in their Lemma 9.1 (for the proof see Lemma 9.2 in Goldberg et al. (2000)). Fiat et al. (2002) claim that the result of Goldberg et al. (2001), which applies to deterministic auctions, also applies to random auctions through a straightforward generalization. We focus on deterministic auctions, which is without loss of generality under assumptions of regularity.

such that for given $(\mathbf{v}, \mathbf{c})$, *the number of trades is the highest index* $k \in \{0, 1, ..., \min\{n, m\}\}$ *satisfying, in the case of two-sided private information,* $e_k(v_{(k)}, c_{[k]}; \mathbf{v}_{(k+1)}, \mathbf{c}_{[k+1]}) \geq 0$, *and in the case of one-sided private information pertaining to buyers* $e_k^B(v_{(k)}; \mathbf{v}_{(k+1)}, \mathbf{c}) \geq 0$ *and sellers* $e_k^S(c_{[k]}; \mathbf{v}, \mathbf{c}_{[k+1]}) \geq 0$, *where payments are threshold payments.*

Proof. See Appendix B.

Lemma 1 provides a characterization of the data on which the estimated virtual type functions can depend, subject to incentive compatibility, individual rationality, and envy freeness. It tells us that for estimating whether the k-th best trader should trade, it is without loss of generality to use only data from traders who are less efficient. These data are exactly the information one can elicit from agents who exit in a clock auction. Thus, although Proposition 1 shows that in two-sided settings clock auctions are a restriction relative to Bayesian optimal mechanisms, for the purposes of estimation, clock auctions are not restrictive. Clock auctions are without loss of generality for implementing prior-free mechanisms that are BOAEE.

Now consider the task of finding consistent estimators of the weighted virtual type functions. The following lemma, which relates the expected inverse hazard rates to expected spacings, with the expectations being taken with respect to the true distributions, tells us that the task of finding consistent estimators for the weighted virtual type functions boils down to finding consistent estimators of the spacings between order statistics for the buyers' values and for the sellers' costs.

Lemma 2 $jE_{\mathbf{v}|F,...,F}[v_{(j)} - v_{(j+1)}] = E_{\mathbf{v}|F,...,F}\left[\frac{1-F(v_{(j)})}{f(v_{(j)})}\right]$ *and* $jE_{\mathbf{c}|G,...,G}[c_{[j+1]} - c_{[j]}] = E_{\mathbf{c}|G,...,G}\left[\frac{G(c_{[j]})}{g(c_{[j]})}\right]$.

Proof. See Appendix B.

Lemma 2 suggests that the weighted virtual value function, evaluated at the j-th highest value, $\Phi_\alpha(v_{(j)}) = v_{(j)} - \alpha\frac{1-F(v_{(j)})}{f(v_{(j)})}$, can be estimated by $v_{(j)} - \alpha j \sigma_j^v$, where σ_j^v is an estimate of the expected spacing between $v_{(j)}$ and $v_{(j+1)}$. Similarly, it suggests that the weighted virtual cost function, evaluated at the j-th lowest cost, $\Gamma_\alpha(c_{[j]}) = c_{[j]} + \alpha\frac{G(c_{[j]})}{g(c_{[j]})}$, can be estimated by $c_{[j]} + \alpha j \sigma_j^c$, where σ_j^c is an estimate of the expected spacing between $c_{[j+1]}$ and $c_{[j]}$.

Pursuing this idea, we establish the prior-free optimality of the prior-free clock auction $\mathcal{C}_{\phi,\gamma,\tau}$ defined as follows: For any state $\boldsymbol{\omega}$ with buyer clock price p^B, seller clock price p^S, and an equal number $j-1$ of active buyers and sellers (which implies that the values $v_{(j)}, ..., v_{(n)}$ and costs $c_{[j]}, ..., c_{[m]}$ are known from the exit prices of the inactive buyers and sellers and so can be used for estimation), we first define the buyer and seller functions and then define the target function. The buyer and seller functions are

$$\phi(\boldsymbol{\omega}) = p^B - \chi_{\alpha,j}\sigma_j^v \text{ and } \gamma(\boldsymbol{\omega}) = p^S + \chi_{\alpha,j}\sigma_j^c, \tag{7}$$

where $\chi_{\alpha,j}$ is a nonnegative coefficient and σ_j^v and σ_j^c are spacing estimators. For reasons that will be clear when we discuss our consistency requirement, for $j \in \{1, ..., \min\{m, n\}\}$, we let

$$\chi_{\alpha,j} \equiv \max\{0,\ \alpha(j-2) - (1-\alpha)\}. \tag{8}$$

It follows that $\chi_{\alpha,j}$ is nonnegative (ensuring that the mechanism is deficit free) and has the property that $\operatorname{plim}_{i\to\infty} \chi_{\alpha,i}/i = \alpha$ (so that $\chi_{\alpha,j}$ has the asymptotic properties of αj). To achieve consistent estimators of the expected spacing between values $v_{(j-1)}$ and $v_{(j)}$ and between the costs $c_{[j]}$ and $c_{[j-1]}$, derived based on values $v_{(j)}, ..., v_{(n)}$ and costs $c_{[j]}, ..., c_{[m]}$, we use the average of r_n and r_m spacings for nearby worse types. Specifically, given exit prices for buyers of $\hat{v}_{(n)}, ..., \hat{v}_{(j)}$ and exit prices for sellers of $\hat{c}_{[m]}, ..., \hat{c}_{[j]}$ (equal to the corresponding true values and costs under agents' dominant strategies), we let

$$\sigma_j^v \equiv \begin{cases} \frac{\hat{v}_{(j)} - \hat{v}_{(j+\min\{r_n, n-j\})}}{\min\{r_n, n-j\}}, & \text{if } j < n \\ \frac{1}{n+1}, & \text{otherwise} \end{cases} \text{ and } \sigma_j^c \equiv \begin{cases} \frac{\hat{c}_{[j+\min\{r_m, m-j\}]} - \hat{c}_{[j]}}{\min\{r_m, m-j\}}, & \text{if } j < m \\ \frac{1}{m+1}, & \text{otherwise,} \end{cases} \tag{9}$$

where r_j satisfies

$$\lim_{j\to\infty} r_j = \infty \quad \text{and} \quad \lim_{j\to\infty} \frac{r_j}{j} = 0. \tag{10}$$

With the assumptions of continuity and that agents play their dominant strategies, (10) ensures consistent estimation of the spacings. In particular, letting $\lfloor x \rceil$ denote x rounded to the nearest integer, then given $\rho \in (0, 1)$, (10) ensures that $\sigma_{\lfloor \rho n \rceil}^v$ converges in probability to $E_{\mathbf{v}}[v_{(\lfloor \rho n \rceil)} - v_{(\lfloor \rho n \rceil + 1)}]$ as n grows large and $\sigma_{\lfloor \rho m \rceil}^c$ to $E_{\mathbf{c}}[c_{[\lfloor \rho m \rceil + 1]} - c_{[\lfloor \rho m \rceil]}]$ as m grows large. (We discuss and illustrate the rate of convergence in the online appendix.)

It remains to specify a target function. For the purposes of asymptotic optimality, the target function is irrelevant because the specification of a target only affects the number of trades by at most one and so does not affect the asymptotic properties of the mechanism. A suitable choice for the target estimator can be defined analogously to that of the BOCA. For a state $\boldsymbol{\omega}$ with $j-1$ active buyers and sellers, we estimate $\tau(\boldsymbol{\omega}) \in [\phi(\boldsymbol{\omega}), \gamma(\boldsymbol{\omega})]$ that maximizes the probability of $j-1$ trades when there should be $j-1$ trades, i.e., maximizes the probability that $v_{(j-1)} - \chi_{\alpha,j-1}\sigma_{j-1}^v \geq \tau(\boldsymbol{\omega}) \geq c_{[j-1]} + \chi_{\alpha,j-1}\sigma_{j-1}^c$ when $v_{(j-1)} - \chi_{\alpha,j-1}\sigma_{j-1}^v \geq c_{[j-1]} + \chi_{\alpha,j-1}\sigma_{j-1}^c$, under some assumptions on the distribution of $v_{(j-1)}$ and $c_{[j-1]}$. For example, one could use the following target estimator:[36]

$$\tau(\boldsymbol{\omega}) = \min\left\{\gamma(\boldsymbol{\omega}),\ \max\left\{\phi(\boldsymbol{\omega}),\ \frac{\phi(\boldsymbol{\omega})+\gamma(\boldsymbol{\omega})}{2} + \left(1-\frac{\alpha}{2}\right)\left(\sigma_j^v - \sigma_j^c\right)\right\}\right\}. \tag{11}$$

According to (23), the target virtual type is the midpoint between $\phi(\boldsymbol{\omega})$ and $\gamma(\boldsymbol{\omega})$ plus $(1-\alpha/2)(\sigma_j^v - \sigma_j^c)$. The second term moves the target upward (closer to $\gamma(\boldsymbol{\omega})$) if $\sigma_j^v > \sigma_j^c$ to account for the expectation that an exit on the seller side is more likely than on the buyer side for equal movements in the virtual types. Conversely, the adjustment is made in the opposite direction if $\sigma_j^v < \sigma_j^c$. The adjustment is greater the lower is α, reflecting the increased value of avoiding an exit when the weight on efficiency is larger. As mentioned above, McAfee (1992) uses the midpoint between the standing clock prices as the target, which corresponds to the midpoint between $v_{(j)}$ and $c_{[j]}$.

We can now prove the following result.

Proposition 3 *Assuming the regularity condition holds, the prior-free clock auction defined by* (19)–(23) *is prior-free optimal.*

The result that our prior-free clock auction is BOAEE follows because its structure mirrors that of the BOCA. To prove the asymptotic optimality portion of Proposition 3 (which technically only requires the uniqueness of the prices satisfying (3) and the Bayesian optimality of price posting in the large), we begin by showing in Lemma 3 that uniform bounds exist for the variance, denoted $V[\cdot]$, of $\chi_{\alpha,j}\sigma_j^v$ and $\chi_{\alpha,j}\sigma_j^c$ away from the boundary. Next, Lemma 4 shows that the difference between the theoretical and prior-free virtual types is uniformly convergent in probability to zero away from the boundary. Lemmas 3 and 4 consider the buyer and seller sides of the market separately and so simply consider limits as m and n go to infinity. Lemma 5 then combines the two sides of the market, considering η-fold replicas.[37] Lemma 5 uses the preceding lemmas to show that fixing m and n and considering η-fold replicas of the economy, as η goes to infinity, the share of agents who trade in a clock auction based on estimated virtual values $\tilde{\Phi}_\alpha(j) \equiv v_{(j)} - \chi_{\alpha,j}\sigma_j^v$ and estimated virtual costs $\tilde{\Gamma}_\alpha(j) \equiv c_{[j]} + \chi_{\alpha,j}\sigma_j^c$ approaches the share in the optimal mechanism. Intuitively, if $\tilde{\Phi}_\alpha$ and $\tilde{\Gamma}_\alpha$ stay close to Φ_α and Γ_α, then the first intersection point of $\tilde{\Phi}_\alpha$ and $\tilde{\Gamma}_\alpha$ cannot be far from the (unique) intersection of Φ_α and Γ_α. Proposition 3 then follows from the fact that in both the optimal mechanism and the prior-free clock auction, it is the highest-valuing buyers and lowest-cost sellers who trade, and that in both cases the payments are in an interval bounded

[36] This estimator maximizes the probability that $v_{(j-1)} - \chi_{\alpha,j-1}\sigma_{j-1}^v \geq \tau(\boldsymbol{\omega}) \geq c_{[j-1]} + \chi_{\alpha,j-1}\sigma_{j-1}^c$ when using σ_j^v as an estimate of σ_{j-1}^v, and σ_j^c for σ_{j-1}^c, and assuming that $v_{(j-1)}$ is distributed uniformly between $v_{(j)}$ and $v_{(j)} + 2\sigma_j^v$ and similarly for $c_{[j-1]}$ between $c_{[j]} - 2\sigma_j^c$ and $c_{[j]}$.

[37] If we did not let both the supply and demand side grow proportionally, asymptotic optimality would be rather trivial. For example, with private information only on the buyers' side and a fixed supply of k units (with $c_k < \overline{v}$), the Vickrey auction would be asymptotically optimal.

by the trading agent with the worst type and the nontrading agent with the best type and so differ by at most one spacing.

Lemma 3 *Given $\overline{\rho} \in (0,1)$, there exist $u^v(\rho,n)$ and $u^c(\rho,m)$ that are increasing in ρ and converge to zero as n and m increase to infinity such that for all n and m sufficiently large and all $\rho \in [0,\overline{\rho}]$, $V\left[\chi_{\alpha,\lfloor\rho n\rceil}\sigma^v_{\lfloor\rho n\rceil}\right] \leq u^v(\rho,n)$ and $V\left[\chi_{\alpha,\lfloor\rho m\rceil}\sigma^c_{\lfloor\rho m\rceil}\right] \leq u^c(\rho,m)$.*

Proof. See Appendix B.

Given Lemma 3, we can prove uniform convergence in probability of the theoretical and smoothed virtual types.

Lemma 4 *Given $\overline{\rho} \in (0,1)$, the difference between the theoretical and smoothed virtual values $\Phi_\alpha(v_{(\lfloor\rho n\rceil)}) - \tilde{\Phi}_\alpha(\lfloor\rho n\rceil)$ on $(\overline{v} - \overline{\rho}(\overline{v} - \underline{v}), \overline{v})$ is uniformly convergent in probability to zero, and the difference between the theoretical and smoothed virtual costs $\Gamma_\alpha(c_{[\lfloor\rho m\rceil]}) - \tilde{\Gamma}_\alpha(\lfloor\rho m\rceil)$ on $(\underline{c}, \underline{c} + \overline{\rho}(\overline{c} - \underline{c}))$ is uniformly convergent in probability to zero.*

Proof. See Appendix B.

Given Lemma 4, we can now show that the number of trades in the prior-free clock auction approaches that in the optimal mechanism.

Lemma 5 *Assuming the uniqueness of the prices satisfying* (3) *and the Bayesian optimality of price posting in the large, given m and n and considering η-fold replicas of the economy, as η goes to infinity, the share of agents who trade in the prior-free clock auction defined by* (19)–(23) *converges in probability to the share in the optimal mechanism.*

Proof. See Appendix B.

The adaptation of Proposition 3 to the case of one-sided private information is straightforward. In that case, the virtual type need not be estimated on the side of the market without private information, and the target virtual type need not be estimated at all. For example, with private information only on the buyer side, if the state $\boldsymbol{\omega}$ has $j-1$ active buyers, then $\gamma(\boldsymbol{\omega})$ in the analysis above would be replaced simply by $c_{[j]}$, and $\tau(\boldsymbol{\omega})$ would be replaced by $c_{[j-1]}$.

Sequential consistency

As shown in Proposition 3, the prior-free clock auction defined in (19)–(23) is prior-free optimal. As we now show, it also satisfies an additional property, which we define in this section and which relates to the sequential consistency of a dynamic mechanism. To define sequential consistency, we now assume that the designer delegates the operation of a dynamic mechanism to a decision maker, such as an auctioneer. We consider whether the auctioneer's following the protocol defined by the mechanism is credible in the sense of maximizing the auctioneer's expected payoff. This is similar to the credibility notion developed independently by Akbarpour and Li (2017). Sequential consistency differs from their definition of a credible extensive form game plus strategy profile by specifying how the auctioneer forms expectations. It addresses the commitment problem faced by an auctioneer regarding when to stop a clock auction, assuming that the auctioneer can commit to running a clock auction.[38]

Suppose that, at time t, a prior-free clock auction generates an estimate $\hat{\mu}_t$ observable to the auctioneer of a distribution over a vector of payoff-relevant types $\boldsymbol{\theta}$. Let A be the set of actions available to the auctioneer and let $u(a,\boldsymbol{\theta})$ be the auctioneer's payoff when taking action a and the types are $\boldsymbol{\theta}$. We say that a plan of action (a mapping from histories into A) that prescribes the action a_t after a history that generates estimate $\hat{\mu}_t$ over $\boldsymbol{\theta}$ is *sequentially consistent with respect to estimate* $\hat{\mu}_t$ *of* $\boldsymbol{\theta}$ if $a_t \in \arg\max_{a\in A} E_{\boldsymbol{\theta}|\hat{\mu}_t}[u(a,\boldsymbol{\theta})]$. That is, the prescribed action a_t maximizes the auctioneer's expected payoff when his expectation is taken using the estimate $\hat{\mu}_t$ generated by the mechanism.

To apply this notion to the auctioneer's decision whether to stop a clock auction, let Ω_j be the set of clock auction states that follow a buyer or seller exit that results in $j-1$ active buyers and $j-1$ active sellers. We assume that agents follow their obviously dominant strategies of bidding truthfully so that the observed exits recorded by the states in Ω_j reveal $\mathbf{v}_{(j)}$ and $\mathbf{c}_{[j]}$. Assume that the auctioneer's payoff is $-\infty$ if the clock auction ends with a deficit and is otherwise equal to, or perfectly aligned with, the designer's objective that puts weight α on revenue. We assume that the auctioneer's action set is such that following an exit that results in equal numbers of active buyers and sellers, the auctioneer can choose whether to end the auction or continue (with commitment to then end the auction should the target prices be achieved without exits).

We consider whether given $\boldsymbol{\omega} \in \Omega_j$, the plan by the auctioneer to stop the auction if and only if $\phi(\boldsymbol{\omega}) \geq \gamma(\boldsymbol{\omega})$ is sequentially consistent with respect to beliefs generated by σ_j^v and σ_j^c that $v_{(j-1)} = v_{(j)} + \sigma_j^v$ and $c_{[j-1]} = c_{[j]} - \sigma_j^c$. For this to hold, at each state $\boldsymbol{\omega} \in \Omega_j$ in the clock auction with clock prices p^B and p^S, it must be that $\phi(\boldsymbol{\omega}) \geq \gamma(\boldsymbol{\omega})$ if and only if (i) $p^B \geq p^S$ (to ensure no deficit) and (ii) $E_{v_{(j-1)}-v_{(j)}|\sigma_j^v}\left[\Phi_\alpha(v_{(j-1)}) \mid \mathbf{v}_{(j)}\right] \geq E_{c_{[j]}-c_{[j-1]}|\sigma_j^c}\left[\Gamma_\alpha(c_{[j-1]}) \mid \mathbf{c}_{[j]}\right]$.

We can show that the prior-free clock auction defined in (19)–(23) satisfies this condition. To see this, note that we can write (ii) as

$$v_{(j)} + \sigma_j^v - \alpha E_{v_{(j-1)}-v_{(j)}|\sigma_j^v}\left[\frac{1-F(v_{(j-1)})}{f(v_{(j-1)})} \mid \mathbf{v}_{(j)}\right] \geq c_{[j]} - \sigma_j^c + \alpha E_{c_{[j]}-c_{[j-1]}|\sigma_j^c}\left[\frac{G(c_{[j-1]})}{g(c_{[j-1]})} \mid \mathbf{c}_{[j]}\right],$$

[38]McAdams and Schwarz (2007) analyze a setup in which not even that level of commitment is possible, finding a role for delay costs, reputation, and intermediaries.

which using Lemma 2, we can rewrite as

$$v_{(j)} + \sigma_j^v - \alpha(j-1)\sigma_j^v \geq c_{[j]} - \sigma_j^c + \alpha(j-1)\sigma_j^c$$

or, rearranging, as

$$v_{(j)} - (\alpha(j-2) - (1-\alpha))\sigma_j^v \geq c_{[j]} + (\alpha(j-2) - (1-\alpha))\sigma_j^c. \tag{12}$$

Because we are considering a state in which the clock prices are defined by exits, $p^B = v_{(j)}$ and $p^S = c_{[j]}$. Thus, both $p^B \geq p^S$ and (12) hold if and only if

$$p^B - \chi_{\alpha,j}\sigma_j^v \geq p^S + \chi_{\alpha,j}\sigma_j^c, \tag{13}$$

where $\chi_{\alpha,j}$ is defined in (20). Because (13) is the criterion for ending the auction in the prior-free optimal clock auction defined by (19)–(23), the auctioneer's incentives are aligned with the auction protocol, and so sequential consistency is satisfied. Thus, we have the following proposition:

Proposition 4 *Assuming the regularity condition holds, the stopping rule for the auctioneer in the prior-free optimal clock auction defined in* (19)–(23) *is sequentially consistent with respect to the beliefs generated by the spacing estimators* $\boldsymbol{\sigma}^v$ *and* $\boldsymbol{\sigma}^c$*. Furthermore, it is the unique such clock auction up to the definition of the spacing estimators* $\boldsymbol{\sigma}^v$ *and* $\boldsymbol{\sigma}^c$ *and the target function.*

To illustrate, consider a two-sided setting with $\alpha = 1$. When $j-1$ buyers and sellers remain active in the clock auction and the estimated spacings are σ_j^v and σ_j^c, then the estimated increase in revenue from continuing until there is an additional exit on both sides is the additional revenue of $\sigma_j^v + \sigma_j^c$ from the remaining $j-2$ trading pairs, less the revenue $v_{(j)} - c_{[j]}$ from the one lost trade:

$$(j-2)(\sigma_j^v + \sigma_j^c) - (v_{(j)} - c_{[j]}).$$

The estimated loss in social surplus from continuing is the estimated surplus from the lost trade, $v_{(j)} - c_{[j]} + \sigma_j^v + \sigma_j^c$. A confidence interval for the estimated term $\sigma_j^v + \sigma_j^c$ can be constructed using a bootstrap approach (see, e.g., Silverman, 1986, Chapter 6.4).[39] This is illustrated in Figure 2 for an example with 20 buyers and 20 sellers.[40] Panel (a) shows

[39]Given $\mathbf{v}_{(j)}$, generate a bootstrap sample by taking a uniform random selection of $n-j+1$ elements of $\mathbf{v}_{(j)}$ with replacement and adjusting those with error terms drawn from the uniform kernel and calibrating to the mean and variance of $\mathbf{v}_{(j)}$. This bootstrap sample then implies a bootstrap value for the spacing estimator. Repeating this procedure allows one to construct a bootstrap confidence interval.

[40]As an alternative, for certain choices of r_n and r_m, asymptotic normality results can be used to derive confidence bounds by noting the equivalence between our estimator and the hazard rate estimator based on

the estimated 95% confidence bands for the estimated virtual types assuming $\alpha = 1$, given the data available following the exit of the j-th highest valuing buyer and j-th, and panel (b) shows the corresponding confidence bands for the estimated increase in revenue from continuing the auction.

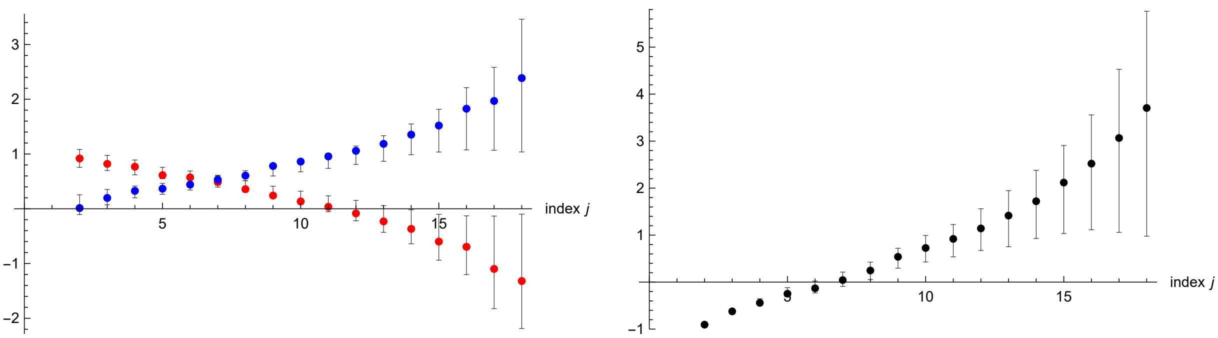

(a) Estimated virtual types with CIs

(b) Estimated revenue changes with CIs

Figure 2: Panel (a): Bootstrap 95% confidence bounds for estimated virtual types with $\alpha = 1$ given $\mathbf{v}_{(j)}$ and $\mathbf{c}_{[j]}$. Panel (b): Bootstrap 95% confidence bounds for the increase in revenue from continuing the clock auction until an additional buyer and seller exit following the exit of the j^{th} highest valuing buyer and lowest valuing seller. Results assume $n = 20$ and $r_n = n^{4/5}$, with values and costs drawn from the uniform distribution on [0,1].

As can be seen from Figure 2(a), the estimated virtual value function first exceeds the estimated virtual cost function following the exit of the 6-th highest-value buyer and 6-th lowest-cost seller (at index $j = 6$). Thus, according to the rules of our prior-free optimal clock auction (setting aside the target function for purposes of the illustration), the auction would end following this exit. Turning to Figure 2(b), following the exit of the 7-th highest-value buyer and 7-th lowest-cost seller (at index $j = 7$), the expected revenue change from continuing the auction remains positive, but following the exit of the 6-th best agents, the expected change is negative. Thus, the auctioneer's dynamic incentive is to continue the auction until after the exit of the 6-th best agents and then to end the auction, consistent with the protocol defined by the mechanism.

Criteria for selecting the estimator

Although as a matter of statistics there is some flexibility in fixing the precise details of the estimators used in a prior-free optimal clock auction, given Lemma 2, sequential consistency

the empirical cdf and nearest neighbor density estimator (see footnote 41). On the asymptotic normality of the empirical cdf, see van der Vaart (1998, p. 165). As shown by Moore and Yackel (1977, Theorem 2), the nearest neighbor density estimator is asymptotically normal when $r_n = n^{2/3}$ and f has bounded first derivative. Thus, attaining asymptotic normality requires faster convergence of $\frac{r_n}{n}$ to zero than with $r_n = n^{4/5}$, which as described below is required (up to proportionality) for minimizing mean square error.

requires that the virtual type estimators be based on expected spacings between types. That suggests the use of spacing estimators, such as $\boldsymbol{\sigma}^v$ and $\boldsymbol{\sigma}^c$, that are from the class of nearest neighbor estimators, which are estimators based on the average of nearby spacings.[41] Nearest neighbor estimators have bias and variance that go to zero with n and m when r_j satisfies (10), and their mean square error is minimized when r_j is proportional to $j^{4/5}$.[42] Thus, the nearest neighbor estimators that attain the minimum mean square error are essentially unique in that they are uniquely defined up to proportionality constants.

We summarize with the following proposition.

Proposition 5 *Assuming the regularity condition holds, the prior-free clock auction defined in* (19)–(23)*, with $r_j \propto j^{4/5}$, is the unique prior-free optimal clock auction that is sequentially consistent and uses a virtual type estimator that achieves the minimum mean square error among nearest neighbor estimators, up to proportionality and the definition of the target estimator.*

5 Extensions

Thus far, we have focused on a designer whose objective is the weighted sum of revenue and social surplus. However, our design is flexible enough to incorporate a variety of alternative objectives and additional constraints. In particular, we can incorporate any constraint that can be stated in terms of adjustments to the functions defining a clock auction. Here we briefly comment on how a designer facing heterogenous groups of agents can implement caps on the number of units that a particular group can buy or sell, impose minimal revenue requirements in order for trade by a group to occur, or favor certain groups over others.

[41]Following Silverman (1986, Chapter 2.5), the estimator $j\sigma_j^v$ of $\frac{1-F(v_{(j)})}{f(v_{(j)})}$ can be derived from the empirical cdf and a "nearest neighbor" density estimator, which coincides with the kernel density estimator based on the uniform kernel and bandwidth given data $\mathbf{v}_{(j)}$ of $h(\mathbf{v}_{(j)}) = r_n\sigma_j^v$, and similarly for the estimator $j\sigma_j^c$ of $\frac{G(c_{[j]})}{g(c_{[j]})}$. More specifically, given kernel $\bar{k}(x)$, for $-1 < x < 1$, with bandwidth $\bar{h}(\mathbf{x})$, the unbounded kernel density estimator given data $\mathbf{x}$ is $\kappa_{\mathbf{x}}(y) \equiv \frac{1}{h(\mathbf{x})|\mathbf{x}|}\sum_{i=1}^{|\mathbf{x}|}\bar{k}\left(\frac{y-x_{(i)}}{h(\mathbf{x})}\right)$. The kernel density estimator given data $\mathbf{x}$ and upper bound $\max_i x_i$ is calculated by reflecting the data across the upper bound and then truncating the resulting density, i.e., using

$$\kappa_{\mathbf{x}}^{ub}(y) \equiv 2\kappa_{(x_{(\dim \mathbf{x})},\ldots,x_{(2)},x_{(1)},x_{(1)},2x_{(1)}-x_{(2)},\ldots,2x_{(1)}-x_{(\dim \mathbf{x})})}(y),$$

for $y \leq \max_i x_i$. If we use the uniform kernel $\bar{k}(x) \equiv \frac{1}{2}1_{0<|x|\leq 1}$ with bandwidth $\bar{h}(\mathbf{x}) \equiv x_{(1)} - x_{(1+r_n)}$, which implies $\bar{h}(\mathbf{v}_{(j)}) = r_n\sigma_j^v$, the kernel density estimate for f at $v_{(j)}$ given data $\mathbf{v}_{(j)}$ is $\bar{f}(\mathbf{v}_{(j)}) \equiv \frac{n-j+1}{n+1}\kappa_{\mathbf{v}_{(j)}}^{ub}(v_{(j)}) = \frac{1}{n+1}\frac{1}{\sigma_j^v}$, and so the estimated inverse hazard rate is $\left(1 - \frac{n-j+1}{n+1}\right)/\bar{f}(\mathbf{v}_{(j)}) = j\sigma_j^v$.

[42]An r_j-nearest neighbor estimator has mean square error of order $\left(\frac{r_j}{j}\right)^4 + \frac{1}{r_j}$, which is minimized when r_j is proportional to $j^{4/5}$ (Silverman, 1986, Chapters 3 and 5.2.2), in which case the approximate value of the mean integrated square error tends to zero at the rate $j^{-4/5}$ (Silverman, 1986, Chapter 3.7.2). For an illustration of the rate convergence, see the online appendix.

For the purposes of this section, we assume that agents have characteristics that are observable to the designer, so that the designer can a priori place subsets of agents into groups of symmetric agents while allowing for asymmetries across different groups. For example, traders of carbon emission permits might be identifiable as either power plants, cement manufacturers, or other manufacturers, with traders within a group being symmetric, but with the possibility of asymmetries across groups.

In the online appendix, we show how our analysis and clock auction generalize to such settings. As described in the online appendix, we refer to a clock auction that accommodates differences across groups as a *discriminatory clock auction*. It differs from a nondiscriminatory clock auction in that, loosely speaking, there are separate but coordinated clocks for each buyer and seller group.

Group-specific quantity caps

A designer or regulator may want to cap the number of units that a subset of buyers acquires. More generally, constraints of this kind can be described by a partition matroid as in Dütting et al. (2017), which allows the feasible trading set to be defined by a maximum number of agents from each of different buyer groups and seller groups. Such constraints can be imposed within a discriminatory clock auction by treating that subset of buyers for which there is a cap as a group and starting the procedure by advancing the clock price for that group until the number of active agents in the group is reduced to the number eligible to trade. This is implemented by defining the discriminatory clock auction mappings so that the stopping rule cannot be satisfied until the cap for the group is met.

Revenue constraints

A discriminatory clock auction can also accommodate the requirement that members of some buyer group $\hat{b}$ contribute payments of at least $\underline{R}$ in order for any members of that group to trade. This is accomplished by setting the estimated virtual value for group $\hat{b}$ buyers equal to minus infinity as long as that group's clock price times the number of active buyers in the group remains below $\underline{R}$.

Favoring groups of agents

A discriminatory clock auction is also flexible enough to allow the designer to favor a particular subset of agents over others.[43] This can be accomplished using a discriminatory clock auction that assigns favored and non-favored agents to separate groups, evaluates non-favored agents using virtual types that incorporate the designer's unconstrained weight α

[43] For example, in the case of U.S. federal acquisitions, the "Buy American Act" specifies favoritism for domestic bidders and domestic small business bidders. (U.S. Federal Acquisition Regulation, FAR 25.105(b))

on revenue, and evaluates favored agents using virtual types with weight α^f on revenue. Favoritism then simply means $\alpha^f < \alpha$.

6 Conclusions

We develop a prior-free clock auction that is Bayesian optimal absent estimation error and asymptotically optimal. As a clock auction, it endows agents with obviously dominant strategies to bid truthfully and preserves the privacy of trading agents. Methodologically, we exploit the connection between the empirical measure of spacings between order statistics and the theoretical construct of virtual types.

Many features of the mechanisms we develop, such as Bayesian optimality absent estimation error and the flexibility to accommodate various constraints and to pursue a combination of revenue and surplus goals, may prove useful in various setups and applications. While our setup is general in that it accommodates situations in which private information pertains to one or both sides of the market, in many markets traders decide endogenously whether they act as buyers or as sellers. Extending the methodology of the present paper to account for the endogeneity of traders' positions—buy, sell, or hold—seems a promising avenue for future research.

The prior-free mechanism design approach raises the somewhat philosophical question as to why a designer who is not endowed with a prior should be interested in asymptotic Bayesian optimality in the first place. A possible answer to this question is that asymptotic optimality provides a reassuring evaluation and consistency criterion. Asked how well his mechanism performs, a designer employing an asymptotically optimal mechanism facing many traders may find it reassuring to know that he would not have chosen any other mechanism at the outset had he then known the distributions that he has inferred now. In that sense, asymptotic optimality, like privacy preservation, protects the designer from regret.

A Appendix: Definition of a two-sided clock auction

The formal definition of a two-sided clock auction is given below. Following the definition, we comment on the adaptation required for a one-sided setup.

Clock auction $\mathcal{C}_{\phi,\gamma,\tau}$ is defined as follows: For $t \in \{0, 1, ...\}$, the state of a clock auction is $\boldsymbol{\omega}_t = (z_t, \boldsymbol{\omega}_t^B, \boldsymbol{\omega}_t^S)$ where $z_t \in \{0, 1\}$ specifies whether the clock auction has ended ($z_t = 1$) or not ($z_t = 0$), and $\boldsymbol{\omega}_t^B = (\mathbb{N}^A, \mathbf{x}^B, p^B)$ and $\boldsymbol{\omega}_t^S = (\mathbb{M}^A, \mathbf{x}^S, p^S)$ are buyer and seller states. The components of the buyer state are: the set of active buyers $\mathbb{N}^A \subseteq \mathbb{N}$ with cardinality n^A, the vector of exit prices for non-active buyers $\mathbf{x}^B \in \mathcal{R}^{n-n^A}$, and the buyer clock price $p^B \in \mathcal{R}$. The seller state has an analogous structure. Let Ω be the set of all possible states.

The clock auction starts in state $\boldsymbol{\omega}_0 \equiv (0, \boldsymbol{\omega}_0^B, \boldsymbol{\omega}_0^S)$, where $\boldsymbol{\omega}_0^B = (\mathbb{N}, \emptyset, \underline{p})$ and $\boldsymbol{\omega}_0^S = (\mathbb{M}, \emptyset, \overline{p})$ with $\underline{p} < \underline{v}$ and $\overline{p} > \overline{c}$, so that initially all agents are active. The clock auction continues until a state is reached that has a first component equal to 1, at which point the active buyers and sellers trade, with the active buyers paying the buyer clock price and the active sellers receiving the seller clock price.

We require that $\phi : \Omega \to \mathcal{R}$ is increasing in p^B, that $\gamma : \Omega \to \mathcal{R}$ is increasing in p^S, and that $\tau : \Omega \to \mathcal{R}$ satisfies $\tau(\boldsymbol{\omega}_t) \in [\phi(\boldsymbol{\omega}_t), \gamma(\boldsymbol{\omega}_t)]$ whenever $\phi(\boldsymbol{\omega}_t) \leq \gamma(\boldsymbol{\omega}_t)$. Define target buyer price $T^B(\boldsymbol{\omega}_t)$ to be the buyer clock price such that $\phi(\boldsymbol{\omega}'_t)$ is equal to $\tau(\boldsymbol{\omega}_t)$, where $\boldsymbol{\omega}'_t = (z_t, (\mathbb{N}^A, \mathbf{x}^B, T^B(\boldsymbol{\omega}_t)), \boldsymbol{\omega}_t^S)$, i.e., $\boldsymbol{\omega}'_t$ is equal to $\boldsymbol{\omega}_t$ but with the buyer clock price p^B replaced by the target buyer price $T^B(\boldsymbol{\omega}_t)$. Similarly define target seller price $T^S(\boldsymbol{\omega}_t)$ to be the value for the seller clock price that equates $\gamma(\boldsymbol{\omega}''_t)$ with $\tau(\boldsymbol{\omega}_t)$, where $\boldsymbol{\omega}''_t = (z_t, \boldsymbol{\omega}_t^B, (\mathbb{M}^A, \mathbf{x}^S, T^S(\boldsymbol{\omega}_t)))$.

For $t \in \{0, 1, ...\}$, if $\boldsymbol{\omega}_t^B = (\mathbb{N}^A, \mathbf{x}^B, p^B)$, $\boldsymbol{\omega}_t^S = (\mathbb{M}^A, \mathbf{x}^S, p^S)$, and $z_t = 0$, state $\boldsymbol{\omega}_{t+1}$ is determined as follows:

If $n^A = m^A$: If $n^A = 0$ or $\phi(\boldsymbol{\omega}_t) \geq \gamma(\boldsymbol{\omega}_t)$, then $\boldsymbol{\omega}_{t+1} = (1, \boldsymbol{\omega}_t^B, \boldsymbol{\omega}_t^S)$. Otherwise, proceed as follows (the choice of which clock price to adjust first is arbitrary; clock prices can also be adjusted simultaneously): Increase the buyer clock price from p^B until either a buyer i exits at clock price $\hat{p}^B$, in which case $\boldsymbol{\omega}_{t+1}^B = (\mathbb{N}^A \backslash \{i\}, (\mathbf{x}^B, \hat{p}^B), \hat{p}^B)$, or the buyer clock price reaches $T^B(\boldsymbol{\omega}_t)$ with no exit, in which case $\boldsymbol{\omega}_{t+1}^B = (\mathbb{N}^A, \mathbf{x}^B, T^B(\boldsymbol{\omega}_t))$. Decrease the seller clock price from p^S until either a seller j exits at $\hat{p}^S$, in which case $\boldsymbol{\omega}_{t+1}^S = (\mathbb{M}^A \backslash \{j\}, (\mathbf{x}^S, \hat{p}^S), \hat{p}^S)$, or the seller clock price reaches $T^S(\boldsymbol{\omega}_t)$ with no exit, in which case $\boldsymbol{\omega}_{t+1}^S = (\mathbb{M}^A, \mathbf{x}^S, T^S(\boldsymbol{\omega}_t))$. If both target prices are reached with no exits, then $z_{t+1} = 1$; otherwise $z_{t+1} = 0$.

If $n^A > m^A$, increase the buyer clock price from p^B until either a buyer i exits at $\hat{p}^B$, in which case $\boldsymbol{\omega}_{t+1}^B = (\mathbb{N}^A \backslash \{i\}, (\mathbf{x}^B, \hat{p}^B), \hat{p}^B)$, $\boldsymbol{\omega}_{t+1}^S = \boldsymbol{\omega}_t^S$, and $z_{t+1} = 0$, or the buyer clock price reaches $\overline{p}$, in which case $\boldsymbol{\omega}_{t+1}^B = (\hat{\mathbb{N}}^A, \mathbf{x}^B, \overline{p})$, where $\hat{\mathbb{N}}^A$ consists of m^A randomly selected elements of $\mathbb{N}^A$, $\boldsymbol{\omega}_{t+1}^S = \boldsymbol{\omega}_t^S$, and $z_{t+1} = 1$.

If $n^A < m^A$, decrease the seller clock price from p^S until a seller j exits at $\hat{p}^S$, in which case $\boldsymbol{\omega}_{t+1}^B = \boldsymbol{\omega}_t^B$, $\boldsymbol{\omega}_{t+1}^S = (\mathbb{M}^A \backslash \{j\}, (\mathbf{x}^S, \hat{p}^S), \hat{p}^S)$, and $z_{t+1} = 0$, or the seller clock price reaches

$\underline{p}$, in which case $\boldsymbol{\omega}_{t+1}^S = (\hat{\mathbb{M}}^A, \mathbf{x}^S, \underline{p})$, where $\hat{\mathbb{M}}^A$ consists of n^A randomly selected elements of $\mathbb{M}^A$, $\boldsymbol{\omega}_{t+1}^B = \boldsymbol{\omega}_t^B$, and $z_{t+1} = 1$.

The above definition of a clock auction for the two-sided setup is easily adapted to the case of one-sided private information on the buyer side by, essentially, eliminating the seller clock and the seller state. To be precise, for state $\boldsymbol{\omega}_t$ with n^A active buyers, let $\tau(\boldsymbol{\omega}_t) \equiv c_{[n^A]}$ and define the corresponding target buyer price $T^B(\boldsymbol{\omega}_t)$ as above. Given $z_t = 0$ and $\boldsymbol{\omega}_t^B = (\mathbb{N}^A, \mathbf{x}^B, p^B)$ with $n^A > 0$, $\boldsymbol{\omega}_{t+1}$ is determined as follows: If $\phi(\boldsymbol{\omega}_t) \geq c_{[n^A+1]}$, then $\boldsymbol{\omega}_{t+1} = (1, \boldsymbol{\omega}_t^B, \boldsymbol{\omega}_t^S)$. Otherwise, increase the buyer clock price until either a buyer i exits at clock price $\hat{p}^B$, in which case $\boldsymbol{\omega}_{t+1}^B = (\mathbb{N}^A \backslash \{i\}, (\mathbf{x}^B, \hat{p}^B), \hat{p}^B)$ and $z_{t+1} = 0$, or the buyer clock price reaches $T^B(\boldsymbol{\omega}_t)$ with no exit, in which case $\boldsymbol{\omega}_{t+1}^B = (\mathbb{N}^A, \mathbf{x}^B, T^B(\boldsymbol{\omega}_t))$ and $z_{t+1} = 1$. Symmetric adjustments are made for the case of one-sided information on the seller side.

B Appendix: Proofs

Proof of Proposition 2. The data available to the designer with $j-1$ buyers and $j-1$ sellers active in a clock auction are $\mathbf{v}_{(j)}$ and $\mathbf{c}_{[j]}$. Given $\mathbf{v}_{(j)}$, $\mathbf{c}_{[j]}$, and $\Phi_\alpha(v_{(j)}) < \Gamma_\alpha(c_{[j]})$ for some $j \in \{2, ..., \min\{m, n\} + 1\}$, maximizing the probability that, if $\Phi_\alpha(v_{(j-1)}) \geq \Gamma_\alpha(c_{[j-1]})$, the quantity traded is $j-1$ boils down to choosing the target virtual type $\delta_j \in [\Phi_\alpha(v_{(j)}), \Gamma_\alpha(c_{[j]})]$ that maximizes the probability that $\Gamma_\alpha(c_{[j-1]}) \leq \delta_j \leq \Phi_\alpha(v_{(j-1)})$. This probability is given by

$$P(\delta) \equiv \left(1 - F_{(j-1)}\left(\Phi_\alpha^{-1}(\delta)\right)\right) G_{[j-1]}\left(\Gamma_\alpha^{-1}(\delta)\right) = \left(\frac{(1 - F\left(\Phi_\alpha^{-1}(\delta)\right))G\left(\Gamma_\alpha^{-1}(\delta)\right)}{(1 - F(v_{(j)}))G(c_{[j]})}\right)^{j-1},$$

where the equality follows because $1 - F_{(j-1)}(v) = \left(\frac{1-F(v)}{1-F(v_{(j)})}\right)^{j-1}$ for $v > v_{(j)}$ and $G_{[j-1]}(c) = \left(\frac{G(c)}{G(c_{[j]})}\right)^{j-1}$ for $c < c_{[j]}$. Let $\delta^* \in \arg\max_{\delta \in [\underline{c}, \overline{v}]} P(\delta)$. If $P(\delta)$ is quasiconcave,[44] δ^* is given by the first-order condition

$$P'(\delta^*) = P(\delta^*)(j-1)\left[-\frac{f(\Phi_\alpha^{-1}(\delta^*))}{1 - F(\Phi_\alpha^{-1}(\delta^*))}\frac{1}{\Phi_\alpha'(\Phi_\alpha^{-1}(\delta^*))} + \frac{g(\Gamma_\alpha^{-1}(\delta^*))}{G(\Gamma_\alpha^{-1}(\delta^*))}\frac{1}{\Gamma_\alpha'(\Gamma_\alpha^{-1}(\delta^*))}\right] = 0,$$

which completes the proof. ■

Proof of Lemma 1. The proof focuses on the two-sided setup, the logic for the one-sided settings following along similar lines. Envy freeness and ex post individual rationality imply

[44] A sufficient condition for quasiconcavity is that the hazard rate $\lambda_F = f/(1-F)$ is increasing and concave and the hazard rate $\lambda_G = g/G$ decreasing and concave. To see this, notice that $P' = P(j-1)[-\lambda_F^3/(\lambda_F^2 + \alpha\lambda_F') + \lambda_G^3/(\lambda_G^2 + \alpha\lambda_G')]$. The sign of P'' at δ such that $P'(\delta) = 0$ is the same as that of the derivative of $[-\lambda_F^3/(\lambda_F^2 + \alpha\lambda_F') + \lambda_G^3/(\lambda_G^2 + \alpha\lambda_G')]$, which is negative under the stated conditions.

that when there are k trades, it is the k highest valuing buyers and lowest cost sellers who trade and that trading buyers pay the same $p_k^B \in [v_{(k+1)}, v_{(k)}]$ and trading sellers receive the same $p_k^S \in [c_{[k]}, c_{[k+1]}]$. Dominant strategy incentive compatibility implies it must be possible to determine whether k agents trade irrespective of the types $\mathbf{v}_{(k-1)}$ and $\mathbf{c}_{[k-1]}$, giving rise to the existence of the evaluation functions. Further, dominant strategy incentive compatibility implies that p_k^B and p_k^S are the worst types that a agent could report and still trade, i.e., threshold payments. Because these worst types must be the same for all trading agents, they cannot depend on the types of trading agents on the same side of the market. ■

Proof of Lemma 2. Take the case of costs. The proof for values is analogous. For $j \in \{1, ..., m-1\}$, the density of the j-th lowest order statistic out of m draws from distribution G is $\frac{m!}{(j-1)!(m-j)!} G^{j-1}(x)(1-G(x))^{m-j} g(x)$. It then follows that

$$\begin{aligned}
E_{\mathbf{c}}\left[\frac{G(c_{[j]})}{g(c_{[j]})}\right] &= \int_{\underline{c}}^{\overline{c}} \frac{G(x)}{g(x)} \frac{m!}{(j-1)!(m-j)!} G^{j-1}(x)(1-G(x))^{m-j} g(x) dx \\
&= \int_{\underline{c}}^{\overline{c}} \frac{m!}{(j-1)!(m-j)!} G^{j}(x)(1-G(x))^{m-j} dx \\
&= (m-j) \int_{\underline{c}}^{\overline{c}} \frac{m!}{(j-1)!(m-j)!} x G^{j}(x)(1-G(x))^{m-j-1} g(x) dx \\
&\quad -j \int_{\underline{c}}^{\overline{c}} \frac{m!}{(j-1)!(m-j)!} x G^{j-1}(x)(1-G(x))^{m-j} g(x) dx \\
&= j \int_{\underline{c}}^{\overline{c}} \frac{m!}{j!(m-j-1)!} x G^{j}(x)(1-G(x))^{m-j-1} g(x) dx \\
&\quad -j \int_{\underline{c}}^{\overline{c}} \frac{m!}{(j-1)!(m-j)!} x G^{j-1}(x)(1-G(x))^{m-j} g(x) dx \\
&= j E_{\mathbf{c}}[c_{[j+1]} - c_{[j]}],
\end{aligned}$$

where the first equality uses the definition of the expectation, the second rearranges, the third uses integration by parts, the fourth rearranges, and the fifth again uses the definition of the expectation. ■

Proof of Lemma 3. We show the result for costs. The result for values follows analogously. In this proof, where we use notation such as ρm as an index, e.g., $c_{[\rho m]}$, we mean that ρm is rounded to the nearest integer. Recall that $\chi_{\alpha,j} \equiv \max\{0,\ \alpha(j-2)-(1-\alpha)\}$. Because $\chi_{0,\rho m} = 0$, the result holds for $\alpha = 0$, so assume $\alpha > 0$. Because $\overline{\rho} < 1$, we are away from the boundary and can focus on $\sigma_{\rho m}^c = \frac{c_{[\rho m + r_m]} - c_{[\rho m]}}{r_m}$ ($\sigma_{\rho m}^c$ is defined differently close to the boundary, i.e., for $\rho m > m - r_m$). To show that $V(\chi_{\alpha,\rho m} \sigma_{\rho m}^c)$ goes to zero with m, it is sufficient (given assumptions of continuity on a bounded support) to show that

$\chi^2_{\alpha,\rho m} V\left[\frac{G(c_{[\rho m+r_m]})-G(c_{[\rho m]})}{r_m}\right]$ goes to zero, and because the cdf of an order statistic is itself a uniform order statistic, it is sufficient to show that $\chi^2_{\alpha,\rho m} V\left[\frac{u_{[\rho m+r_m]}-u_{[\rho m]}}{r_m}\right]$ goes to zero, where $u_{[i]}$ is the i-th order statistic out of m draws from $U[0,1]$.

Results for uniform order statistics imply that $V\left[u_{[i]}\right] = \frac{i(m+1-i)}{(m+1)^2(m+2)}$ and $Cov\left[u_{[i]}u_{[j]}\right] = \frac{i(m+1-j)}{(m+1)^2(m+2)}$, so for $i<j$, we have

$$
\begin{aligned}
V\left(\frac{u_{[i+r_m]}-u_{[i]}}{r_m}\right) &= \frac{1}{r_m^2}V\left(u_{[i+r_m]}-u_{[i]}\right) = \frac{1}{r_m^2}\left(V(u_{[i+r_m]})+V(u_{[i]}\right)-2Cov\left(u_{[i]}u_{[i+r_m]}\right) \\
&= \frac{m+1-r_m}{(m+1)^2(m+2)r_m}.
\end{aligned}
$$

It follows that $\chi^2_{\alpha,\rho m} V\left[\frac{u_{[\rho m+r_m]}-u_{[\rho m]}}{r_m}\right] = \chi^2_{\alpha,\rho m}\frac{m+1-r_m}{(m+1)^2(m+2)r_m}$, which is non-decreasing in ρ. Taking the limit of the above expression, we have $\lim_{m\to\infty}\chi^2_{\alpha,\rho m}\frac{m+1-r_m}{(m+1)^2(m+2)r_m} = \lim_{m\to\infty}\frac{\chi^2_{\alpha,\rho m}}{m^2}\frac{1}{r_m} = 0$, where the first equality uses $\lim_{m\to\infty}\frac{r_m}{m}=0$ and the second equality uses $\lim_{m\to\infty} r_m = \infty$ and the fact that $\chi_{\alpha,\rho m}$ is of order m. This gives us the existence of a uniform bound. ■

Proof of Lemma 4. Again, we show the result for virtual costs, with the result for virtual values following analogously. Also as above, where we use ρm as an index, we mean that ρm is rounded to the nearest integer. Two auxiliary results will be useful. First, one can show that for $\rho\in(0,1)$,

$$\lim_{m\to\infty}\rho m E_{\mathbf{c}}\left[c_{[\rho m+1]}-c_{[\rho m]}-\frac{c_{[\rho m+r_m]}-c_{[\rho m]}}{r_m}\right]=0. \tag{14}$$

To see this, note that given $\rho\in(0,1)$ and m sufficiently large, $\rho m+r_m\le m$, so the expression in (14) is well defined. We can then write the expression inside the limit in (14) as

$$
\begin{aligned}
& \rho m E_{\mathbf{c}}\left[c_{[\rho m+1]}-c_{[\rho m]}-\frac{c_{[\rho m+r_m]}-c_{[\rho m]}}{r_m}\right] \\
= \; & \rho m E_{\mathbf{c}}\left[c_{[\rho m+1]}-c_{[\rho m]}-\sum_{i=1}^{r_m}\frac{c_{[\rho m+i]}-c_{[\rho m+i-1]}}{r_m}\right] \\
= \; & \rho m E_{\mathbf{c}}\left[\frac{1}{\rho m}\frac{G(c_{[\rho m]})}{g(c_{[\rho m]})}-\sum_{i=1}^{r_m}\frac{1}{(\rho m+i-1)r_m}\frac{G(c_{[\rho m+i-1]})}{g(c_{[\rho m+i-1]})}\right] \\
= \; & \rho m E_{\mathbf{u}}\left[\frac{1}{\rho m}\frac{u_{[\rho m]}}{g(G^{-1}(u_{[\rho m]}))}-\sum_{i=1}^{r_m}\frac{1}{(\rho m+i-1)r_m}\frac{u_{[\rho m+i-1]}}{g(G^{-1}(u_{[\rho m+i-1]}))}\right] \\
= \; & E_{\mathbf{u}}\left[\frac{u_{[\rho m]}}{g(G^{-1}(u_{[\rho m]}))}-\sum_{i=1}^{r_m}\frac{\frac{\rho m}{r_m}}{\rho m+i-1}\frac{u_{[\rho m+i-1]}}{g(G^{-1}(u_{[\rho m+i-1]}))}\right] \\
= \; & \frac{\frac{\rho m}{m+1}}{g(G^{-1}(\frac{\rho m}{m+1}))}-\sum_{i=1}^{r_m}\frac{\frac{\rho m}{r_m}}{\rho m+i-1}\frac{\frac{\rho m+i-1}{m+1}}{g(G^{-1}(\frac{\rho m+i-1}{m+1}))},
\end{aligned}
$$

where the first equality writes $c_{[\rho m+r_m]}-c_{[\rho m]}$ as the sum of r_m spacings, the second equality uses Lemma 2, the third equality uses the fact that a cdf evaluated at an order statistic is itself a uniform order statistic, with $u_{[j]}$ denoting the j-th order statistic out of m draws from $U[0,1]$, the fourth equality rearranges, and the fifth equality uses $E_{\mathbf{u}}\left[u_{[j]}\right]=\frac{j}{m+1}$. Taking the limit of the above expression as m goes to infinity, we get

$$
\begin{aligned}
& \lim_{m\to\infty}\rho m E_{\mathbf{c}}\left[c_{[\rho m+1]}-c_{[\rho m]}-\frac{c_{[\rho m+r_m]}-c_{[\rho m]}}{r_m}\right] \\
= \; & \lim_{m\to\infty}\frac{\frac{\rho m}{m+1}}{g(G^{-1}(\frac{\rho m}{m+1}))}-\sum_{i=1}^{r_m}\frac{\frac{\rho m}{r_m}}{\rho m+i-1}\frac{\frac{\rho m+i-1}{m+1}}{g(G^{-1}(\frac{\rho m+i-1}{m+1}))} \\
= \; & \frac{\rho}{g(G^{-1}(\rho))}-\frac{\rho}{g(G^{-1}(\rho))}\lim_{m\to\infty}\sum_{i=1}^{r_m}\frac{\frac{\rho m}{r_m}}{\rho m+i-1}=0,
\end{aligned}
$$

where the final equality follows from the fact that $\lim_{m\to\infty}\sum_{i=1}^{r_m}\frac{\frac{\rho m}{r_m}}{\rho m+i-1}=1$. (To see this, note that $\sum_{i=1}^{r_m}\frac{\frac{\rho m}{r_m}}{\rho m+i-1}\leq 1$ and that $\sum_{i=1}^{r_m}\frac{\frac{\rho m}{r_m}}{\rho m+i-1}\geq\frac{\rho}{\rho+\frac{r_m}{m}-\frac{1}{m}}\to_{m\to\infty}1$, where the limit uses $\lim_{m\to\infty}\frac{r_m}{m}=0$ and $\lim_{m\to\infty}\frac{1}{m}=0$.)

Second, given $\overline{\rho}\in(0,1)$,

$$
\lim_{m\to\infty}\sup_{\rho\in[0,\overline{\rho}]}V_{\mathbf{c}}\left[\frac{G(c_{[\rho m]})}{g(c_{[\rho m]})}\right]=0. \tag{15}
$$

To see this, let U_j be the distribution of the j-th lowest order statistic out of m draws from $U[0,1]$. Using the fact that the cdf of an order statistic is a uniform order statistic with the notation above, we have

$$\begin{aligned} V_{\mathbf{c}}\left[\frac{G(c_{[\rho m]})}{g(c_{[\rho m]})}\right] &= V_{\mathbf{u}}\left[\frac{u_{[\rho m]}}{g\big(G^{-1}(u_{[\rho m]})\big)}\right] = E_{\mathbf{u}}\left[\left(\frac{u_{[\rho m]}}{g\big(G^{-1}(u_{[\rho m]})\big)}\right)^2\right] - \left(\frac{\frac{\rho m}{m+1}}{g\big(G^{-1}\big(\frac{\rho m}{m+1}\big)\big)}\right)^2 \\ &= \int_0^1 \left(\frac{x}{g(G^{-1}(x))}\right)^2 dU_{\rho m}(x) - \left(\frac{\frac{\rho m}{m+1}}{g\big(G^{-1}\big(\frac{\rho m}{m+1}\big)\big)}\right)^2. \end{aligned}$$

Taking the limit, we have

$$\lim_{m\to\infty} V_{\mathbf{c}}\left[\frac{G(c_{[\rho m]})}{g(c_{[\rho m]})}\right] = \lim_{m\to\infty} \int_0^1 \left(\frac{x}{g(G^{-1}(x))}\right)^2 dU_{\rho m}(x) - \left(\frac{\rho}{g(G^{-1}(\rho))}\right)^2 = 0,$$

where the final equality follows because $(\frac{x}{g(G^{-1}(x))})^2$ is bounded on $[0,1]$ and, in the limit as m goes to infinity, $U_{\rho m}$ places point mass on ρ.

Using these two auxiliary results along with Lemma 2, we now proceed with the proof of Lemma 4. Denote the difference between theoretical and smoothed virtual costs as

$$Y_j \equiv \Gamma_\alpha(c_{[j]}) - \tilde{\Gamma}_\alpha(j) = \alpha\frac{G(c_{[j]})}{g(c_{[j]})} - \chi_{\alpha,j}\sigma_j^c.$$

For j away from the boundary, $\sigma_j^c = \frac{c_{[j+r_m]}-c_{[j]}}{r_m}$, so $Y_j = \alpha\frac{G(c_{[j]})}{g(c_{[j]})} - \chi_{\alpha,j}\frac{c_{[j+r_m]}-c_{[j]}}{r_m}$ and

$$\begin{aligned} E_{\mathbf{c}}\left[Y_j\right] &= \alpha E_{\mathbf{c}}\left[\frac{G(c_{[j]})}{g(c_{[j]})}\right] - \chi_{\alpha,j}E_{\mathbf{c}}\left[\frac{c_{[j+r_m]}-c_{[j]}}{r_m}\right] \\ &= \alpha j E_{\mathbf{c}}[c_{[j+1]} - c_{[j]}] - \chi_{\alpha,j}E_{\mathbf{c}}\left[\frac{c_{[j+r_m]}-c_{[j]}}{r_m}\right], \end{aligned}$$

where the final equality uses Lemma 2.

Given $\overline{\rho} \in (0,1)$, for all $\rho \in (0,\overline{\rho})$, there exists $\overline{m}$ sufficiently large such that for all $m > \overline{m}$, $\rho m + r_m \leq m$ (implying that the expression above is well defined and $\sigma_{\rho m}^c = \frac{c_{[\rho m+r_m]}-c_{[\rho m]}}{r_m}$) and $\chi_{\alpha,\rho m} = \alpha(\rho m - 2) - (1-\alpha)$. It follows that

$$\begin{aligned} E_{\mathbf{c}}\left[Y_{\rho m}\right] &= \alpha\rho m E_{\mathbf{c}}[c_{[\rho m+1]} - c_{[\rho m]}] - \left(\alpha(\rho m - 2) - (1-\alpha)\right)E_{\mathbf{c}}\left[\frac{c_{[\rho m+r_m]}-c_{[\rho m]}}{r_m}\right] \\ &= \alpha\rho m E_{\mathbf{c}}\left[c_{[\rho m+1]} - c_{[\rho m]} - \frac{c_{[\rho m+r_m]}-c_{[\rho m]}}{r_m}\right] + (1+\alpha)E_{\mathbf{c}}\left[\frac{c_{[\rho m+r_m]}-c_{[\rho m]}}{r_m}\right]. \end{aligned}$$

Taking the limit, we have

$$\lim_{m\to\infty} E_{\mathbf{c}}\left[Y_{\rho m}\right] = \lim_{m\to\infty} \alpha\rho m E_{\mathbf{c}}\left[c_{[\rho m+1]} - c_{[\rho m]} - \frac{c_{[\rho m+r_m]} - c_{[\rho m]}}{r_m}\right] = 0,$$

where the first equality uses $\lim_{m\to\infty} E_{\mathbf{c}}\left[\frac{c_{[\rho m+r_m]}-c_{[\rho m]}}{r_m}\right] = 0$ (the numerator is bounded by $\overline{c} - \underline{c}$ and r_m goes to infinity with m) and the second equality uses (14).

Turning to the variance of $Y_{\rho m}$, we can write $V_{\mathbf{c}}[Y_{\rho m}]$ as

$$V_{\mathbf{c}}[Y_{\rho m}] = \alpha^2 V_{\mathbf{c}}\left[\frac{G(c_{[\rho m]})}{g(c_{[\rho m]})}\right] + V_{\mathbf{c}}\left[\chi_{\alpha,\rho m}\sigma^c_{\rho m}\right] + \text{covariance term.}$$

By (15), $\lim_{m\to\infty}\sup_{\rho\in[0,\overline{\rho}]} V_{\mathbf{c}}\left[\frac{G(c_{[\rho m]})}{g(c_{[\rho m]})}\right] = 0$. By Lemma 3, for m sufficiently large, $V_{\mathbf{c}}[\chi_{\alpha,\rho m}\sigma^c_{\rho m}] \leq u^c(\rho,m) \leq u^c(\overline{\rho},m)$, where $\lim_{m\to\infty} u^c(\overline{\rho},m) = 0$. Thus, $\lim_{m\to\infty}\sup_{\rho\in[0,\overline{\rho}]} V_{\mathbf{c}}\left[\chi_{\alpha,\rho m}\sigma^c_{\rho m}\right] = 0$. It follows then by the Cauchy-Schwarz inequality, that the limit of the covariance terms is also zero. Thus,

$$\lim_{m\to\infty}\sup_{\rho\in[0,\overline{\rho}]} V_{\mathbf{c}}[Y_{\rho m}] = 0. \tag{16}$$

Using Markov's Theorem, for all $\varepsilon > 0$,

$$\lim_{m\to\infty}\Pr\left(\sup_{\rho\in[0,\overline{\rho}]}|Y_{\rho m}| \geq \varepsilon\right) \leq \lim_{m\to\infty}\sup_{\rho\in[0,\overline{\rho}]}\frac{E_{\mathbf{c}}\left[Y^2_{\rho m}\right]}{\varepsilon^2} = \lim_{m\to\infty}\sup_{\rho\in[0,\overline{\rho}]}\frac{V_{\mathbf{c}}[Y_{\rho m}]}{\varepsilon^2} = 0,$$

where the first equality uses $\lim_{m\to\infty} E_{\mathbf{c}}[Y_{\rho m}] = 0$ and the second equality uses (16). This establishes uniform convergence in probability to zero. ■

Proof of Lemma 5. Fix m and n and consider η-fold replicas of the economy. By the assumption of unique prices satisfying (3) and the Bayesian optimality of price posting, in the limit as η goes to infinity, the optimal mechanism is a price posting mechanism such that buyers are charged the price v^* and trading sellers are paid c^*, with cutoff types v^* and c^* defined by $\Phi_\alpha(v^*) = \Gamma_\alpha(c^*)$ and $n(1-F(v^*)) = mG(c^*)$. For $\alpha = 0$, the result follows from McAfee (1992), so consider the case with $\alpha > 0$. Together with our assumption that $\underline{v} \leq \overline{c}$, this implies that $v^* \in (\underline{v},\overline{v})$ and $c^* \in (\underline{c},\overline{c})$. It follows that there exists $\overline{\rho} \in (0,1)$ sufficiently large such that $v^* \in (\overline{v} - \overline{\rho}(\overline{v}-\underline{v}),\overline{v})$ and $c^* \in (\underline{c},\underline{c}+\overline{\rho}(\overline{c}-\underline{c}))$, which allows us to use Lemma 4.

For $\varepsilon > 0$ sufficiently small, we define v^ε and c^ε to be the cutoff types if the virtual value function were increased to $\Phi_\alpha + \varepsilon$ and virtual cost function were reduced to $\Gamma_\alpha - \varepsilon$, i.e., $\Phi_\alpha(v^\varepsilon) + \varepsilon = \Gamma_\alpha(c^\varepsilon) - \varepsilon$ and $1 - F(v^\varepsilon) = G(c^\varepsilon)$. We can define $v^{-\varepsilon}$ and $c^{-\varepsilon}$ analogously: $\Phi_\alpha(v^{-\varepsilon}) - \varepsilon = \Gamma_\alpha(c^{-\varepsilon}) + \varepsilon$ and $1 - F(v^{-\varepsilon}) = G(c^{-\varepsilon})$. Because Φ_α and Γ_α are continuous functions on compact supports, and so uniformly continuous, it follows that

$$\lim_{\varepsilon\to 0} v^\varepsilon = v^*,\ \ \lim_{\varepsilon\to 0} v^{-\varepsilon} = v^*,\ \ \lim_{\varepsilon\to 0} c^\varepsilon = c^*,\ \text{and}\ \lim_{\varepsilon\to 0} c^{-\varepsilon} = c^*. \tag{17}$$

This is illustrated in Figure 3 (note that values are decreasing and costs are increasing along the horizontal axis).

By Lemma 4, for all $\varepsilon > 0$, $\lim_{\eta\to\infty}\Pr\left(\sup_{\rho\in[0,\overline{\rho}]}\left|\Gamma_\alpha(c_{[\lfloor\eta\rho m\rceil]}) - \tilde{\Gamma}_\alpha(\lfloor\eta\rho m\rceil)\right| \geq \varepsilon\right) = 0$,

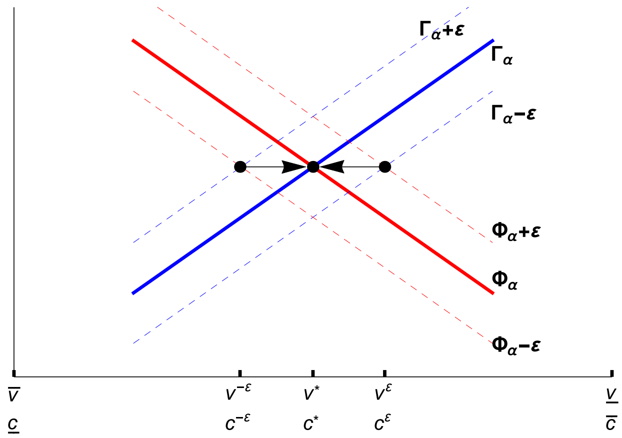

Figure 3: Illustration of $v^\varepsilon < v^* < v^{-\varepsilon}$ and $c^{-\varepsilon} < c^* < c^\varepsilon$ and convergence to v^* and c^*.

which says that in the limit, the probability that for any $\rho \in [0, \overline{\rho}]$, $\tilde{\Gamma}_\alpha(\lfloor \eta\rho m \rceil)$ lies outside of the interval $[\Gamma_\alpha(c_{[\lfloor \eta\rho m \rceil]}) - \varepsilon, \Gamma_\alpha(c_{[\lfloor \eta\rho m \rceil]}) + \varepsilon]$ is zero, and similarly for $\tilde{\Phi}_\alpha(\lfloor \eta\rho n \rceil)$. Thus, in the limit, the probability that the smoothed virtual types intersect at a value $\tilde{v}_{\eta n,\eta m}$ and cost $\tilde{c}_{\eta n,\eta m}$ not bounded by $[v^\varepsilon, v^{-\varepsilon}]$ and $[c^{-\varepsilon}, c^\varepsilon]$ is zero: for all $\varepsilon \in (0, \overline{\varepsilon})$,

$$\lim_{\eta\to\infty} \Pr\left(\tilde{v}_{\eta n,\eta m} \notin [v^\varepsilon, v^{-\varepsilon}]\right) = 0 \text{ and } \lim_{\eta\to\infty} \Pr\left(\tilde{c}_{\eta n,\eta m} \notin [c^{-\varepsilon}, c^\varepsilon]\right) = 0.$$

Thus, using (17), for all $\varepsilon > 0$,

$$\lim_{\eta\to\infty} \Pr\left(|\tilde{c}_{\eta n,\eta m} - c^*| \geq \varepsilon\right) = 0 \text{ and } \lim_{\eta\to\infty} \Pr\left(|v^* - \tilde{v}_{\eta n,\eta m}| \geq \varepsilon\right) = 0,$$

which completes the proof. ■

References

Akbarpour, M. and S. Li (2017): "Credible Mechanism Design," Stanford University.

Arrow, K. (1959): "Toward a Theory of Price Adjustment," in *The Allocation of Economic Resources: Essays in Honor of Bernard Francis Haley*, ed. by M. Abramovitz, Stanford: Stanford University Press, 41–51.

Athey, S. and P. A. Haile (2007): "Nonparametric Approaches to Auctions," in *Handbook of Econometrics*, ed. by J. J. Heckman and E. E. Leamer, Elsevier, vol. 6A, 3847–3966.

Ausubel, L. (2004): "An Efficient Ascending-Bid Auction for Multiple Objects," *American Economic Review*, 94, 1452–1475.

——— (2006): "An Efficient Dynamic Auction for Heterogeneous Commodities," *American Economic Review*, 96, 602–629.

Baliga, S. and R. Vohra (2003): "Market Research and Market Design," *Advances in Theoretical Economics*, 3, Article 5.

Barberà, S., D. Berga, and B. Moreno (2014): "Group Strategy-Proofness in Private Good Economies," SSRN Working Paper 2540583.

Bergemann, D. and S. Morris (2005): "Robust Mechanism Design," *Econometrica*, 73, 1771–1813.

——— (2009): "Robust Implementation in Direct Mechanisms," *Review of Economic Studies*, 76, 1175–1204.

——— (2012): *Robust Mechanism Design*, Hackensack, NJ: World Scientific Publishing.

Bulow, J. and J. Roberts (1989): "The Simple Economics of Optimal Auctions," *Journal of Political Economics*, 97, 1060–1090.

Crémer, J. and R. McLean (1985): "Optimal Selling Strategies Under Uncertainty for a Discriminating Monopolist when Demands are Interdependent," *Econometrica*, 53, 345–361.

——— (1988): "Extraction of the Surplus in Bayesian and Dominant Strategy Auctions," *Econometrica*, 56, 1247–1257.

Cripps, M. W. and J. M. Swinkels (2006): "Efficiency of Large Double Auctions," *Econometrica*, 74, 47–92.

Dasgupta, P. and E. Maskin (2000): "Efficient Auctions," *Quarterly Journal of Economics*, 115, 341–388.

Delacrétaz, D., S. Loertscher, L. M. Marx, and T. Wilkening (2018): "Two-Sided Allocation Problems, Decomposability, and the Impossibility of Efficient Trade," Working Paper, University of Melbourne.

Deshmukh, K., A. V. Goldberg, J. D. Hartline, and A. R. Karlin (2002): "Truthful and Competitive Double Auctions," in *Proceedings of the 10th European Symposium on Algorithms*, Springer-Verlag.

Devanur, N. R., J. D. Hartline, and Q. Yan (2015): "Envy Freedom and Prior-Free Mechanism Design," *Journal of Economic Theory*, 156, 103–143.

Dhangwatnotai, P., T. Roughgarden, and Q. Yan (2015): "Revenue Maximization with a Single Sample," *Games and Economic Behavior*, 91, 318–333.

Dütting, P., T. Roughgarden, and I. Talgam-Cohen (2017): "Modularity and Greed in Double Auctions," *Games and Economic Behavior*, 405, 59–83.

Fiat, A., A. V. Goldberg, J. D. Hartline, and A. R. Karlin (2002): "Competitive Generalized Auctions," in *Conference Proceedings of the Annual ACM Symposium on Theory of Computing*, Montreal, Quebec, Canada: ACM, 72–81.

Gershkov, A., J. K. Goeree, A. Kushnir, B. Moldovanu, and X. Shi (2013): "On the Equivalence of Bayesian and Dominant Strategy Implementation." *Econometrica*, 81, 197–220.

Goldberg, A. V., J. D. Hartline, A. R. Karlin, M. Saks, and A. Wright (2006): "Competitive Auctions," *Games and Economic Behavior*, 55, 242–269.

Goldberg, A. V., J. D. Hartline, and A. Wright (2000): "Competitive Auctions and Digital Goods," Intertrust Starlab Technical Report STAR-TR-99-01.

——— (2001): "Competitive Auctions and Digital Goods," in *Proceedings of the Twelfth Annual ACM-SIAM Symposium on Discrete Algorithms*, ACM/SIAM, 735–744.

Gresik, T. and M. Satterthwaite (1989): "The Rate at which a Simple Market Converges to Efficiency as the Number of Traders Increases: An Asymptotic Result for Optimal Trading Mechanisms," *Journal of Economic Theory*, 48, 304–332.

Guerre, E., I. Perrigne, and Q. Vuong (2000): "Optimal Nonparametric Estimation of First-Price Auctions," *Econometrica*, 68, 525–574.

Hagerty, K. and W. Rogerson (1987): "Robust Trading Mechanisms," *Journal of Economic Theory*, 42, 94–107.

Hurwicz, L. (1973): "The Design of Mechanisms for Resource Allocation," *American Economic Review*, 63, 1–30.

Kagel, J. H. (1995): "Auctions: A Survey of Experimental Research," in *The Handbook of Experimental Economics*, ed. by J. H. Kagel and A. E. Roth, Princeton, NJ: Princeton University Press, 501–585.

Levin, J. and A. Skrzypacz (2016): "Properties of the Combinatorial Clock Auction," *American Economic Review*, 106, 2528–2551.

Li, S. (2017): "Obviously Strategy-Proof Mechanisms," Working Paper, Harvard University.

Loertscher, S., L. Marx, and T. Wilkening (2015): "A Long Way Coming: Designing Centralized Markets with Privately Informed Buyers and Sellers," *Journal of Economic Literature*, 53, 857–897.

Loertscher, S. and C. Mezzetti (2016): "Dominant-Strategy Double-Clock Auctions with Estimation-based Tâtonnement," Working Paper, University of Melbourne.

LOPOMO, G. (1998): "The English Auction Is Optimal among Simple Sequential Auctions," *Journal of Economic Theory*, 82, 144–166.

——— (2000): "Optimality and Robustness of the English Auction," *Games and Economic Behavior*, 36, 219–240.

LUCKING-REILEY, D. (2000): "Vickrey Auctions in Practice: From Nineteenth-Century Philately to Twenty-First-Century E-Commerce," *Journal of Economic Perspectives*, 14, 183–192.

MANELLI, A. M. AND D. R. VINCENT (2010): "Bayesian and Dominant-Strategy Implementation in the Independent Private-Values Model." *Econometrica*, 78, 1905–1938.

MCADAMS, D. AND M. SCHWARZ (2007): "Credible Sales Mechanisms and Intermediaries," *American Economic Review*, 97, 260–276.

MCAFEE, R. P. (1992): "A Dominant Strategy Double Auction," *Journal of Economic Theory*, 56, 434–450.

MCMILLAN, J. (1994): "Selling Spectrum Rights," *Journal of Economic Perspectives*, 8, 145–162.

MILGROM, P. (2004): *Putting Auction Theory to Work*, Cambridge University Press.

MILGROM, P. AND I. SEGAL (2015): "Deferred-Acceptance Auctions and Radio Spectrum Reallocation," Working Paper, Stanford University.

MILGROM, P. AND R. WEBER (1982): "A Theory of Auctions and Competitive Bidding," *Econometrica*, 50, 1089–1122.

MOORE, D. S. AND J. W. YACKEL (1977): "Large Sample Properties of Nearest Neighbor Density Function Estimators," in *Statistical Decision Theory and Related Topics II*, ed. by S. S. Gupta and D. S. Moore, Academic Press, 269–279.

MYERSON, R. AND M. SATTERTHWAITE (1983): "Efficient Mechanisms for Bilateral Trading," *Journal of Economic Theory*, 29, 265–281.

MYERSON, R. B. (1981): "Optimal Auction Design," *Mathematics of Operations Research*, 6, 58–73.

NORMAN, P. (2004): "Efficient Mechanisms for Public Goods with Use Exclusions," *Review of Economic Studies*, 71, 1163–1188.

RUSTICHINI, A., M. A. SATTERTHWAITE, AND S. R. WILLIAMS (1994): "Convergence to Efficiency in a Simple Market with Incomplete Information," *Econometrica*, 62, 1041–1063.

SATTERTHWAITE, M., S. WILLIAMS, AND K. ZACHARIADIS (2015): "Price Discovery Using a Double-Auction," SSRN Working Paper.

SATTERTHWAITE, M. A. AND S. R. WILLIAMS (2002): "The Optimality of a Simple Market Mechanism," *Econometrica*, 70, 1841–1863.

SEGAL, I. (2003): "Optimal Pricing Mechanisms with Unknown Demand," *American Economic Review*, 93, 509–529.

Shapley, L. (1962): "Complements and Substitutes in the Optimal Assignment Problem," *Naval Research Logistics Quarterly*, 9, 45–48.

Silverman, B. W. (1986): *Density Estimation for Statistics and Data Analysis*, Chapman and Hall.

van der Vaart, A. W. (1998): *Asymptotic Statistics*, Cambridge University Press.

Vickrey, W. (1961): "Counterspeculation, Auction, and Competitive Sealed Tenders," *Journal of Finance*, 16, 8–37.

Williams, S. R. (1999): "A Characterization of Efficient, Bayesian Incentive Compatible Mechanisms," *Economic Theory*, 14, 155–180.

Wilson, R. (1987): "Game-Theoretic Analyses of Trading Processes," in *Advances in Economic Theory, Fifth World Congress*, ed. by T. F. Bewley, Cambridge University Press, 33–70.

Online Appendix: Alternative clock auctions

In this appendix we describe two alternative clock auction formats and illustrate the rate of convergence of the optimal prior-free clock auction. In Section 1, we describe discriminatory clock auctions in a setting in which the designer can a priori place subsets of buyers and subsets of sellers into groups of symmetric agents while allowing for asymmetries across different groups. In Section 2, we describe quasi-clock auctions that implement the Bayesian optimal mechanism in a two-sided setup without violating privacy preservation for any trading agents other than the marginal pair. In Section 3, we discuss and illustrate rates of convergence.

1 Discriminatory clock auctions

Generalized Bayesian mechanism design setting

We now allow, without requiring, the possibility that agents have characteristics that are observable to the designer, so that the designer can a priori place subsets of agents into groups of symmetric agents while allowing for asymmetries across different groups. For example, traders of carbon emission permits might be identifiable as either power plants, cement manufacturers, or other manufacturers, with traders within a group being symmetric, but with the possibility of asymmetries across groups.

Let $\mathbb{N}$ and $\mathbb{M}$ denote the sets of buyers and sellers with cardinalities n and m. Let $\mathbb{Z}^B$ and $\mathbb{Z}^S$ be the sets of groups for buyers and sellers when private information pertains to both sides of the market. (When only one side is privately informed, there is, of course, no point distinguishing between groups on the side of the market without private information.) We refer to the setup in which $|\mathbb{Z}^B| = |\mathbb{Z}^S| = 1$ studied thus far as the *symmetric setup*. Let $n^b \geq 1$ be the number of buyers in buyer group b and let $m^s \geq 1$ be the number of sellers in group s, where $n = \sum_{b \in \mathbb{Z}^B} n^b$ and $m = \sum_{s \in \mathbb{Z}^S} m^s$. We assume that at least one buyer group and at least one seller group has 2 or more members. The group membership of each buyer and seller is common knowledge.

Each buyer in group b draws his value independently from the continuously differentiable distribution F^b with support $[\underline{v}^b, \overline{v}^b]$ and positive density f^b, and each seller in group s draws her cost independently from the continuously differentiable distribution G^s with support $[\underline{c}^s, \overline{c}^s]$ and positive density g^s. Each agent is privately informed about his type, but the types and distributions from which they are drawn are unknown to the mechanism designer and the agents. The designer only knows the group identity of each buyer and seller, that agents in the same group draw their types from the same distribution, that group-specific weighted virtual types given by

$$\Phi_\alpha^b(v) \equiv v - \alpha \frac{1 - F^b(v)}{f^b(v)} \quad \text{and} \quad \Gamma_\alpha^s(c) \equiv c + \alpha \frac{G^s(c)}{g^s(c)}$$

are increasing for each buyer and seller group and that the supports satisfy the no-full trade condition

$$\min_{s\in\mathbb{Z}^S}\{\overline{c}^s\} \geq \max_{b\in\mathbb{Z}^B}\overline{v}^b > \min_{s\in\mathbb{Z}^S}\{\underline{c}^s\} \geq \max_{b\in\mathbb{Z}^B}\{\underline{v}^b\},$$

which generalizes condition (2) to the setup with heterogeneous groups.

Under the stipulated assumptions, the allocation rule for the Bayesian optimal mechanism in the setup with two-sided private information can be described as follows: For a given realization of values and costs $(\mathbf{v}, \mathbf{c})$, rank all weighted virtual values in decreasing and all weighted virtual costs in increasing order, irrespective of group membership, and then have all those buyers and sellers trade who would trade in a Walrasian market if weighted virtual values and costs were true values and costs. That is, letting, for all $i \in \mathbb{N}$ in buyer group b and all $j \in \mathbb{M}$ in seller group s, $V_i \equiv \Phi^b_\alpha(v_i)$ and $C_j \equiv \Gamma^s_\alpha(c_j)$ and

$$\mathbf{V} \equiv (V_1, ..., V_n) \quad \text{and} \quad \mathbf{C} \equiv (C_1, ..., C_m),$$

the optimal quantity traded is given by the largest integer k satisfying $V_{(k)} \geq C_{[k]}$, where we use the usual conventions of setting $V_{(0)} = \infty = C_{[m+1]}$ and $C_{[0]} = -\infty = V_{(n+1)}$. In the dominant strategy implementation, trading buyers in group b pay p^b_k and trading sellers in group s receive p^s_k, where

$$p^b_k = \Phi^{b^{-1}}_\alpha\left(\max\{V_{(k+1)}, C_{[k]}\}\right) \quad \text{and} \quad p^s_k = \Gamma^{s^{-1}}_\alpha\left(\min\{C_{[k+1]}, V_{(k)}\}\right).$$

When private information pertains only to buyers, the optimal quantity is the largest index k such that $V_{(k)} \geq c_{[k]}$. In that case, in the dominant strategy implementation, trading buyers in group b pay $p^b_k = \Phi^{b^{-1}}_\alpha\left(\max\{V_{(k+1)}, c_{[k]}\}\right)$. Analogously, when private information pertains only to sellers, the optimal quantity is the largest index k such that $v_{(k)} \geq C_{[k]}$, and in the dominant strategy implementation, trading sellers in group s receive $p^s_k = \Gamma^{s^{-1}}_\alpha\left(\min\{C_{[k+1]}, v_{(k)}\}\right)$.

Discriminatory clock auction

In the generalization of the clock auction to the setup with heterogeneous groups of buyers and sellers, there are separate, but synchronized, clock prices for each buyer group and each seller group. Although buyers in different groups may pay different prices and sellers in different groups may receive different prices, the mechanism remains envy free within groups.

A *discriminatory clock auction* defines state transitions for state space $\hat{\Omega}$, defined below, based on buyer and seller functions $\hat{\phi} : \hat{\Omega} \to \mathcal{R}$ and $\hat{\gamma} : \hat{\Omega} \to \mathcal{R}$ and target function $\hat{\tau} : \hat{\Omega} \to \mathcal{R}$. Thus, we denote a discriminatory clock auction by $\widehat{\mathcal{C}}_{\hat{\phi},\hat{\gamma},\hat{\tau}}$. At $t \in \{0, 1, ...\}$, the state is $\hat{\boldsymbol{\omega}}_t = (z_t, \hat{\boldsymbol{\omega}}^B_t, \hat{\boldsymbol{\omega}}^S_t)$, where $z_t \in \{0, 1\}$ specifies whether the clock auction has ended ($z_t = 1$) or not ($z_t = 0$), $\hat{\boldsymbol{\omega}}^B_t = \times_{b\in\mathbb{Z}^B}\hat{\boldsymbol{\omega}}^b_t$, and $\hat{\boldsymbol{\omega}}^S_t = \times_{s\in\mathbb{Z}^S}\hat{\boldsymbol{\omega}}^s_t$, where $\hat{\boldsymbol{\omega}}^b_t = (\mathbb{N}^{A^b}, \mathbf{x}^b, p^b)$ and

$\hat{\boldsymbol{\omega}}_t^s = (\mathbb{M}^{A^s}, \mathbf{x}^s, p^s)$ are group-specific buyer and seller states with components analogous to the symmetric case. Let $\hat{\Omega}$ be the set of all possible states. We require that $\hat{\phi}$ is increasing in each p^b, that $\hat{\gamma}$ is increasing in each p^s, and that, as in the symmetric case, $\hat{\tau}(\hat{\boldsymbol{\omega}}_t) \in [\hat{\phi}(\hat{\boldsymbol{\omega}}_t), \hat{\gamma}(\hat{\boldsymbol{\omega}}_t)]$ whenever $\hat{\phi}(\hat{\boldsymbol{\omega}}_t) \leq \hat{\gamma}(\hat{\boldsymbol{\omega}}_t)$. The state is initialized as in the symmetric case.

For $t \in \{0, 1, ...\}$, if $z_t = 0$, then $\hat{\boldsymbol{\omega}}_{t+1}$ is determined as follows:

If $\sum_{b \in \mathbb{Z}^B} n^{A^b} = \sum_{s \in \mathbb{Z}^S} m^{A^s}$: If $\sum_{b \in \mathbb{Z}^B} n^{A^b} = 0$ or $\hat{\phi}(\hat{\boldsymbol{\omega}}_t) \geq \hat{\gamma}(\hat{\boldsymbol{\omega}}_t)$, then $\hat{\boldsymbol{\omega}}_{t+1} = (1, \hat{\boldsymbol{\omega}}_t^B, \hat{\boldsymbol{\omega}}_t^S)$. Otherwise, proceed as follows (the order in which clock prices on either side of the market are moved is again immaterial): Increase the vector of buyer clock prices from $\boldsymbol{p}^B = (p^b)_{b \in \mathbb{Z}^B}$ by increasing the clock prices for the smallest number of buyer groups possible so as to increase $\hat{\phi}(\hat{\boldsymbol{\omega}}_t)$ until either there is an exit by a group $\hat{b}$ buyer i at clock price vector $\hat{\mathbf{p}}^B$, in which case $\hat{\boldsymbol{\omega}}_{t+1}^{\hat{b}} = (\mathbb{N}^{A^{\hat{b}}} \backslash \{i\}, (\mathbf{x}^b, \hat{p}^{\hat{b}}), \hat{p}^{\hat{b}})$ and for $b \neq \hat{b}$, $\hat{\boldsymbol{\omega}}_{t+1}^b = (\mathbb{N}^{A^b}, \mathbf{x}^b, \hat{p}^b)$, or the buyer clock prices reach with no exit $\tilde{\mathbf{p}}^B$ such that $\hat{\phi}$ is equal to $\hat{\tau}(\hat{\boldsymbol{\omega}}_t)$ when it is evaluated at the state $\hat{\boldsymbol{\omega}}_t$ with $\boldsymbol{p}^B$ replaced by $\tilde{\mathbf{p}}^B$, in which case for all b, $\hat{\boldsymbol{\omega}}_{t+1}^b = (\mathbb{N}^{A^b}, \mathbf{x}^b, \tilde{p}^b)$. In analogous fashion, decrease the vector of seller clock prices and update the seller state. If both $\hat{\phi}$ and $\hat{\gamma}$, evaluated at the adjusted clock prices, reach the target $\hat{\tau}(\hat{\boldsymbol{\omega}}_t)$ with no exit, then $z_{t+1} = 1$; otherwise $z_{t+1} = 0$.

If $\sum_{b \in \mathbb{Z}^B} n^{A^b} > \sum_{s \in \mathbb{Z}^S} m^{A^s}$, increase only the buyer clock prices as above until there is an exit or all buyer clock prices reach $\overline{p} > \max_{b \in \mathbb{Z}^B} \overline{v}^b$, and similarly for sellers if $\sum_{b \in \mathbb{Z}^B} n^{A^b} < \sum_{s \in \mathbb{Z}^S} m^{A^s}$, with lower bound $\underline{p} < \min_{s \in \mathbb{Z}^S} \underline{c}^s$. The states transition analogously to the symmetric case.

When the auction ends, active buyers pay and active sellers receive their groups' clock prices.

Bayesian optimal discriminatory clock auction

In the one-sided setting, as in the symmetric case, clock auctions are without loss of generality with regard to achieving Bayesian optimality. For example, when only buyers are privately informed and $n > m$, the clock auction first increases the clock prices until there are only m active buyers. Clock prices for the groups are increased so as to maintain the equality of $\Phi_\alpha^b(p^b)$ across the buyer groups that continue to have at least one active agent. If for all buyer groups with at least one active buyer, the clock prices are greater than or equal to the reserve prices $\Phi_\alpha^{b^{-1}}(c_{[m]})$, then the auction ends with each active buyer trading at his group's clock price. Otherwise, the clock auction continues by raising the clock prices for groups that have active buyers in a coordinated fashion until the earlier of the two events—an additional buyer exits or each group b's clock price reaches its target $\Phi_\alpha^{b^{-1}}(c_{[m]})$. If the targets are reached, active buyers trade at these prices. If an exit occurs, target prices update to $\Phi_\alpha^{b^{-1}}(c_{[m-1]})$, and the process continues.

In the two-sided setting, the targets must account for the uncertainty on both sides of the market. Following the exit of the j^b-th highest value buyer in each buyer group b and the exit of the $\hat{\jmath}^s$-th lowest cost seller in each seller group s, and letting $\mathbf{j} \equiv (\times_{b\in\mathbb{Z}^B} j^b, \times_{s\in\mathbb{Z}^S} \hat{\jmath}^s)$, the probability that the remaining active buyers have virtual values greater than or equal to δ and the remaining active sellers have virtual costs less than or equal to δ is

$$\hat{P}_{\mathbf{j}}(\delta) \equiv \prod_{b \text{ s.t. } j^b \geq 2} \left(1 - F^b_{(j^b-1)}\left(\Phi^{b^{-1}}_\alpha(\delta)\right)\right) \prod_{s \text{ s.t. } \hat{\jmath}^s \geq 2} G^s_{[\hat{\jmath}^s-1]}\left(\Gamma^{s^{-1}}_\alpha(\delta)\right).$$

Define $\hat{\delta}_{\mathbf{j}} \in \arg\max_\delta \hat{P}_{\mathbf{j}}(\delta)$. Under quasiconcavity,[45] it is sufficient that $\hat{\delta}_{\mathbf{j}}$ satisfy

$$\sum_{b\in\mathbb{Z}^B_2} (j^b-1) \frac{f^b\left(\Phi^{b^{-1}}_\alpha(\hat{\delta}_{\mathbf{j}})\right)}{1-F^b(\Phi^{b^{-1}}_\alpha(\hat{\delta}_{\mathbf{j}}))} \frac{1}{\Phi^{b\prime}_\alpha\left(\Phi^{b^{-1}}_\alpha(\hat{\delta}_{\mathbf{j}})\right)} = \sum_{s\in\mathbb{Z}^S_2} (\hat{\jmath}^s-1) \frac{g^s\left(\Gamma^{s^{-1}}_\alpha(\hat{\delta}_{\mathbf{j}})\right)}{G^s(\Gamma^{s^{-1}}_\alpha(\hat{\delta}_{\mathbf{j}}))} \frac{1}{\Gamma^{s\prime}_\alpha\left(\Gamma^{s^{-1}}_\alpha(\hat{\delta}_{\mathbf{j}})\right)},$$

where $\mathbb{Z}^B_2$ is the set of buyer groups with $j^b \geq 2$ and $\Phi^b_\alpha(v^b_{(j^b)}) < \hat{\delta}_{\mathbf{j}}$ and $\mathbb{Z}^S_2$ is the set of seller groups satisfying $j^s \geq 2$ and $\Gamma^s_\alpha(c_{[\hat{\jmath}^s]}) > \hat{\delta}_{\mathbf{j}}$.

The following proposition characterizes the *Bayesian optimal discriminatory clock auction (BODCA)*.

Proposition 6 *The BODCA is the discriminatory clock auction $\widehat{\mathcal{C}}_{\hat{\phi},\hat{\gamma},\hat{\tau}}$ that satisfies, for any state $\hat{\boldsymbol{\omega}}_t$ with n^{A^b} active buyers in each buyer group b and m^{A^s} active sellers in each group s,*

$$\hat{\phi}(\hat{\boldsymbol{\omega}}_t) = \min_{b\in\mathbb{Z}^B \text{ s.t. } n^{A^b}\geq 1} \Phi^b_\alpha(v^b_{(n^{A^b}+1)})$$

and

$$\hat{\gamma}(\hat{\boldsymbol{\omega}}_t) = \max_{s\in\mathbb{Z}^S \text{ s.t. } m^{A^s}\geq 1} \Gamma^s_\alpha(c^s_{[m^{A^s}+1]})$$

and

$$\hat{\tau}(\hat{\boldsymbol{\omega}}_t) = \min\left\{\hat{\gamma}(\hat{\boldsymbol{\omega}}_t), \max\left\{\hat{\phi}(\hat{\boldsymbol{\omega}}_t), \hat{\delta}_{(\times_{b\in\mathbb{Z}^B}(n^{A^b}+1), \times_{s\in\mathbb{Z}^S}(m^{A^s}+1))}\right\}\right\}. \tag{18}$$

Proof. Under the BODCA, given the numbers of active buyers and sellers in each group $\mathbf{j} \equiv (j^1, ..., j^{z^B}, \hat{\jmath}^1, ..., \hat{\jmath}^{z^S})$, the optimization problem is to choose δ to maximize

$$\hat{P}(\delta) \equiv \prod_{b\in\hat{\mathbb{Z}}^B \text{ s.t. } j^b\geq 2} \left(1 - F^b_{(j^b)}(\Phi^{b^{-1}}_\alpha(\delta))\right) \prod_{s\in\hat{\mathbb{Z}}^S \text{ s.t. } \hat{\jmath}^s\geq 2} G^s_{[\hat{\jmath}^s]}\left(\Gamma^{s^{-1}}_\alpha(\delta)\right),$$

where $F^b_{(j^b)}$ and $G^s_{[\hat{\jmath}^s]}$ are the distributions of the j^b-th highest value from buyer group b given $v^b_{(j^b+1)}$ and the $\hat{\jmath}^s$-th lowest cost from seller group s given $c^s_{[\hat{\jmath}^s+1]}$. The derivative of $\hat{P}$ can be

[45] Like in the symmetric setting, sufficient conditions for quasiconcavity are that all hazard rates are monotone (increasing for buyers and decreasing for sellers) and concave.

written as

$$\hat{P}'(\delta) = \hat{P}(\delta)\left[-\sum_{b\in\hat{\mathbb{Z}}_2^B} \frac{f^b_{(j^b)}(\Phi_\alpha^{b^{-1}}(\delta))}{1-F^b_{(j^b)}(\Phi_\alpha^{b^{-1}}(\delta))}\Phi_\alpha^{b^{-1}\prime}(\delta) + \sum_{s\in\hat{\mathbb{Z}}_2^S} \frac{g^s_{[\hat{j}^s]}\left(\Gamma_\alpha^{s^{-1}}(\delta)\right)}{G^s_{[\hat{j}^s]}\left(\Gamma_\alpha^{s^{-1}}(\delta)\right)}\Gamma_\alpha^{s^{-1}\prime}(\delta) \right].$$

The first-order condition can thus be written as

$$\sum_{b\in\hat{\mathbb{Z}}_2^B} \frac{f^b_{(j^b)}(\Phi_\alpha^{b^{-1}}(\delta))}{1-F^b_{(j^b)}(\Phi_\alpha^{b^{-1}}(\delta))}\Phi_\alpha^{b^{-1}\prime}(\delta) = \sum_{s\in\hat{\mathbb{Z}}_2^S} \frac{g^s_{[\hat{j}^s]}\left(\Gamma_\alpha^{s^{-1}}(\delta)\right)}{G^s_{[\hat{j}^s]}\left(\Gamma_\alpha^{s^{-1}}(\delta)\right)}\Gamma_\alpha^{s^{-1}\prime}(\delta).$$

Using $1 - F^b_{(j)}(v) = \left(\frac{1-F^b(v)}{1-F^b(v^b_{(j+1)})}\right)^j$ for $v > v^b_{(k+1)}$ and $G^s_{(j)}(c) = \left(\frac{G^s(c)}{G^s(c^s_{[j+1]})}\right)^j$ for $c < c^s_{[j+1]}$, the result follows. ■

BOAEE for discriminatory clock auctions

One way to interpret the Bayesian optimal mechanism with heterogenous groups is, again, in terms of evaluation functions. For example, in the two-sided setting, the evaluation function e_k can be written as $e_k(v_{(k|\mathbf{\Phi})}, c_{[k|\mathbf{\Gamma}]}; \cdot) = V_{(k)} - C_{[k]}$, where $v_{(k|\mathbf{\Phi})}$ is the value associated with the k-th highest (weighted) virtual value $V_{(k)}$ and $c_{[k|\mathbf{\Gamma}]}$ is the cost associated with the k-th lowest virtual cost $C_{[k]}$. In a prior-free setting, the virtual types used for determining $V_{(k)}$ and $C_{[k]}$ have to be estimated. In the presence of heterogenous groups and two-sided private information, a BOAEE mechanism is dominant strategy incentive compatible if and only if the evaluation functions used to determine the allocation are of the form

$$e_k\big(v_{(k|\mathbf{\Phi})}, c_{[k|\mathbf{\Gamma}]}; \mathbf{v}_{(k+1|\mathbf{\Phi})}, \mathbf{c}_{[k+1|\mathbf{\Gamma}]}\big),$$

where $\mathbf{v}_{(k+1|\mathbf{\Phi})}$ is the vector of values associated with the (estimated) $k+1$-st highest and smaller virtual values and $\mathbf{c}_{[k+1|\mathbf{\Gamma}]}$ is the vector of costs associated with the (estimated) $k+1$-st lowest and larger virtual costs. Every trading buyer pays a price equal to the smallest value that he could have reported without affecting the allocation, and every trading seller is paid a price equal to the largest cost he could have reported without affecting the allocation.

It follows that, as in the symmetric setup, for one-sided settings clock auctions are without loss of generality for BOAEE mechanisms. For two-sided settings estimation can be done using clocks without loss of generality, but clock implementation may require sacrificing a trade.

Prior-free optimal discriminatory clock auction

Focusing on the case of two-sided private information, we show how the prior-free optimal clock auction can be generalized to account for differentiated groups of buyers and sellers. The generalization for one-sided private information follows along similar lines.

We define the prior-free discriminatory clock auction $\widehat{\mathcal{C}}_{\hat{\phi},\hat{\gamma},\hat{\tau}}$ that corresponds to the prior-free optimal clock auction $\mathcal{C}_{\phi,\gamma,\tau}$ defined by (19)–(23). For each buyer group b and seller group s, define σ^b_j and σ^s_j to be group-specific spacing estimates, analogous to the symmetric setup. If $\hat{\boldsymbol{\omega}}_t$ indicates, for each buyer group b, a clock price p^b and number of active bidders n^{A^b}, and if $\hat{\boldsymbol{\omega}}_t$ indicates, for each seller group s, a clock price p^s and number of active bidders m^{A^s}, then let

$$\hat{\phi}(\hat{\boldsymbol{\omega}}_t) = \min_{b\in\mathbb{Z}^B \text{ s.t. } n^{A^b}\geq 1} p^b - \chi_{\alpha,n^{A^b}+1}\sigma^b_{n^{A^b}+1} \text{ and } \hat{\gamma}(\hat{\boldsymbol{\omega}}_t) = \max_{s\in\mathbb{Z}^S \text{ s.t. } m^{A^s}\geq 1} p^s + \chi_{\alpha,m^{A^s}+1}\sigma^s_{m^{A^s}+1}.$$

The target function corresponding to (23) is $\hat{\tau}(\hat{\boldsymbol{\omega}}_t) = \min\left\{\hat{\gamma}(\hat{\boldsymbol{\omega}}_t), \max\left\{\hat{\phi}(\hat{\boldsymbol{\omega}}_t), \tilde{\delta}\right\}\right\}$, where $\tilde{\delta}$ satisfies

$$\sum\nolimits_{b\in\tilde{\mathbb{Z}}^B} \frac{1}{\tilde{\delta}-\left(p^b-\chi_{\alpha,n^{A^b}+1}\sigma^b_{n^{A^b}+1}\right)-(2-\alpha)\sigma^b_{n^{A^b}+1}} = \sum\nolimits_{s\in\tilde{\mathbb{Z}}^S} \frac{1}{\left(p^s+\chi_{\alpha,m^{A^s}+1}\sigma^s_{m^{A^s}+1}\right)-\tilde{\delta}-(2-\alpha)\sigma^s_{m^{A^s}+1}},$$

where $\tilde{\mathbb{Z}}^B$ is the set of buyer groups with $n^{A^b} \geq 1$ and $p^b - \chi_{\alpha,n^{A^b}+1}\sigma^b_{n^{A^b}+1} < \tilde{\delta}$ (so that the target function is only defined with respect to buyer groups that still have at least one active buyer and whose estimated virtual types are currently below the target) and $\tilde{\mathbb{Z}}^S$ is the set of seller groups satisfying $m^{A^s} \geq 1$ and $p^s + \chi_{\alpha,m^{A^s}+1}\sigma^s_{m^{A^s}+1} > \tilde{\delta}$.

Using the arguments in the proof of Proposition 3, one can show that the results on asymptotic optimality extend to heterogeneous groups. This combined with arguments analogous to the case of symmetric buyers and symmetric sellers establishes the following result.

Proposition 7 *In the setup with heterogeneous groups of buyers and sellers, there exists a prior-free discriminatory clock auction that is prior-free optimal. Further, the prior-free optimal discriminatory clock auction that is sequentially consistent and attains the minimum mean square error among nearest neighbor estimators is unique up to the virtual target estimator and proportionality constants on parameters of the virtual type estimators.*

2 Quasi-clock auctions

An implication of Proposition 1 is that in two-sided settings the Bayesian optimal mechanism does not permit a clock implementation. The reason for this is that the thresholds for trading on either side of the market—for example, in the symmetric setting $\max\{v_{(k+1)}, \Phi_\alpha^{-1}(\Gamma_\alpha(c_{[k]}))\}$ for buyers and $\min\{c_{[k+1]}, \Gamma_\alpha^{-1}(\Phi_\alpha(v_{(k)}))\}$ for sellers—depend on information that is provided by an agent who optimally trades. A tradeoff thus arises in two-sided settings between the desirable properties of clock auctions and the benefits of Bayesian optimality.

We now briefly discuss how one could augment a clock auction and implement the Bayesian optimal mechanism without violating privacy preservation for any trading agents other than the marginal pair, that is, other than the buyer with value $v_{(k)}$ and the seller with cost $c_{[k]}$ when the optimal quantity traded is k. We refer to the augmented clock auction as a *quasi-clock auction*. To save space, we restrict the discussion to the setting with symmetric buyers and symmetric sellers. The generalization to heterogeneous groups of buyers and sellers is a straightforward extension. Just like the clock auction, a quasi-clock auction consists of two clocks. It proceeds similarly to the clock auction.

Assuming each agent stays active until the clock price equals his type, when the number of active agents on each side of the market is $k-1$ and the buyer and seller clocks stop at the prices $p^B = v_{(k)}$ and $p^S = c_{[k]}$, the buyers with values in the vector $\mathbf{v}_{(k+1)}$ and the sellers with costs in the vector $\mathbf{c}_{[k+1]}$ become inactive as in the clock auction. However, in contrast to the clock auction, the buyer and seller with types $v_{(k)}$ and $c_{[k]}$, who have just exited, may still trade. In particular, they still trade if $\Phi_\alpha(p^B) \geq \Gamma_\alpha(p^S)$, in which case the trading buyers pay $\Phi_\alpha^{-1}(\Gamma_\alpha(p^S))$ and the trading sellers receive $\Gamma_\alpha^{-1}(\Phi_\alpha(p^B))$, rather than the clock prices. If $\Phi_\alpha(p^B) < \Gamma_\alpha(p^S)$, then the quasi-clock auction proceeds until the earlier of the two events: the target prices are reached or an additional agent exits.

In the quasi-clock auction, all agents are price-takers at all times. Consequently, just like the clock auction, the quasi-clock auction endows agents with dominant strategies. It also preserves the privacy of all but at most one trading agent on each side of the market. However, by Li (2017, Theorem 3), it sacrifices the obviousness of the dominant strategies.

Because virtual types can be estimated analogously to the case of clock auctions, prior-free quasi-clock auctions can also be constructed that are BOAEE, asymptotically optimal, and sequentially consistent, and minimize mean square error among nearest neighbor estimators.

The alternative of a quasi-clock auction, which preserves the privacy of almost all trading agents, raises the question as to why one is concerned about privacy preservation. If privacy preservation is desired primarily to protect traders from hold-up by the designer as discussed, e.g., by Lucking-Reiley (2000), then quasi-clock auctions arguably do as good a job as clock auctions, provided buyers and sellers observe the clock prices on the other side of the market, whence they can infer the prices they face. Also, although to a slightly lesser extent than clock auctions, quasi-clock auctions protect the designer from criticism of "money left on the table" because only the marginal traders' values and costs are revealed. Because the difference between the revealed types and prices will typically be "small," quasi-clock auctions will also not perform much worse than clock auctions if the motivation for privacy preservation is post-auction hold-up, e.g., in the form of taxation. Therefore, if quasi-clock auctions do not appear appealing for practical purposes, this may have less to do with their limited ability to preserve privacy than with their failure to satisfy other desiderata such as the obviousness of the dominant strategies and the weak group strategy-proofness this implies.[46]

[46] The findings of Satterthwaite and Williams (2002) that for uniform distributions the efficiency loss of any incentive compatible, ex ante budget balanced mechanism is of the same order as the gains from trade of the marginal pair suggest that departing from clock auctions to always execute this marginal trade may not be worth the cost.

3 Rates of convergence

In this section, we provide simulation results illustrating the performance of a prior-free optimal clock auction in the small. As described in footnote 42, an r_j-nearest neighbor estimator has mean square error of order $\left(\frac{r_j}{j}\right)^4 + \frac{1}{r_j}$, which is minimized when r_j is proportional to $j^{4/5}$ (Silverman, 1986, Chapters 3 and 5.2.2), in which case the approximate value of the mean integrated square error tends to zero at the rate $j^{-4/5}$ (Silverman, 1986, Chapter 3.7.2).

To illustrate performance in the small, we focus on a prior-free optimal clock auction that is sequentially consistent (assuming the regularity condition holds) and uses a virtual type estimator that achieves the minimum mean square error among nearest neighbor estimators for a buyer placing weight $\alpha \in [0,1]$ on revenue. Specifically, we analyze the prior-free optimal clock auction defined as follows: For any state $\boldsymbol{\omega}$ with buyer clock price p^B, seller clock price p^S, and an equal number $j-1$ of active buyers and sellers, the buyer and seller functions are

$$\phi(\boldsymbol{\omega}) = p^B - \chi_{\alpha,j}\sigma_j^v \text{ and } \gamma(\boldsymbol{\omega}) = p^S + \chi_{\alpha,j}\sigma_j^c, \tag{19}$$

where

$$\chi_{\alpha,j} \equiv \max\left\{0,\ \alpha(j-2) - (1-\alpha)\right\}. \tag{20}$$

and

$$\sigma_j^v \equiv \begin{cases} \frac{\hat{v}_{(j)} - \hat{v}_{(j+\min\{r_n, n-j\})}}{\min\{r_n, n-j\}}, & \text{if } j < n \\ \frac{1}{n+1}, & \text{otherwise} \end{cases} \quad \text{and } \sigma_j^c \equiv \begin{cases} \frac{\hat{c}_{[j+\min\{r_m, m-j\}]} - \hat{c}_{[j]}}{\min\{r_m, m-j\}}, & \text{if } j < m \\ \frac{1}{m+1}, & \text{otherwise,} \end{cases} \tag{21}$$

where

$$r_j = j^{4/5}. \tag{22}$$

To define the target function, we initialize the target at $\underline{c}$, i.e., until the state $\boldsymbol{\omega}$ reflects at least one exit on each side of the market, let $\tau(\boldsymbol{\omega}) = \underline{c}$.[47] Once there is at least one exit on each side of the market, if the state shows an equal number $j-1$ of active buyers and sellers, then we define the target function as

$$\tau(\boldsymbol{\omega}) = \min\left\{\gamma(\boldsymbol{\omega}),\ \max\left\{\phi(\boldsymbol{\omega}),\ \frac{\phi(\boldsymbol{\omega}) + \gamma(\boldsymbol{\omega})}{2} + \left(1 - \frac{\alpha}{2}\right)\left(\sigma_j^v - \sigma_j^c\right)\right\}\right\}. \tag{23}$$

As shown in Figure 4(a), this prior-free optimal clock auction achieves over 80% of optimal expected revenue with only six buyer-seller pairs, and it achieves over 75% of optimal social surplus even when there are as few as 2 buyers and 2 sellers. To illustrate results for small markets with α away from the extremes, Figure 4(a) also shows the performance of our

[47] This initialization handicaps the mechanism because it means that there is no possibility that all buyers and sellers could trade. For the comparisons provided here, it seems appropriate to eliminate the possibility that an initial target, necessarily uninformed by any data from the mechanism, happens to permit full trade.

prior-free optimal clock auction relative to the α-optimal mechanism for various intermediate values of α. Generally speaking, the smaller is α, the smaller is the impact of estimation error, and so the better is the performance of the prior-free optimal clock auction. However, the mechanism's use of the first buyer and seller exits for estimation is a greater disadvantage relative to the α-optimal mechanism when n and α are small. Thus, for small numbers of agents, the relative performance of the prior-free optimal clock auction can be better for larger values of α. This is the case in Figure 4(a), where the line for $\alpha = 0$ dips below the lines for $\alpha = 1/4$ and $\alpha = 1/2$ when n is small.

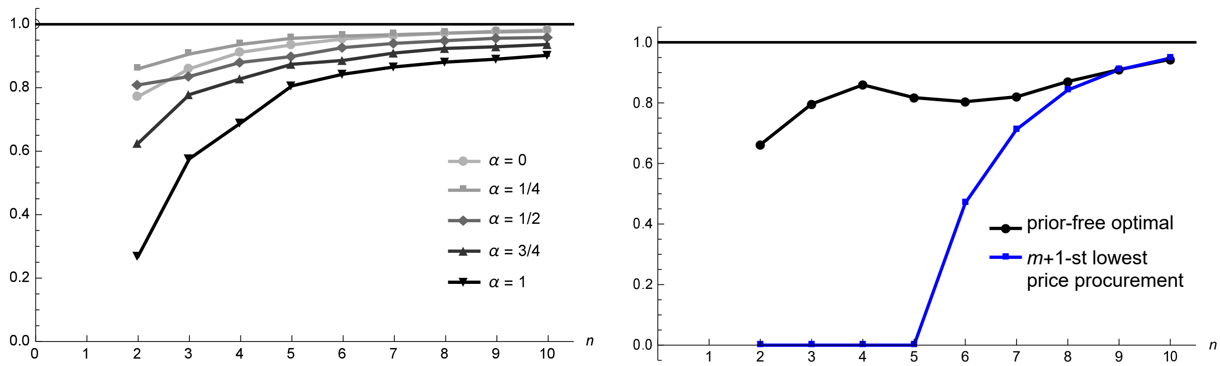

(a) Two-sided: Prior-free performance relative to the α-optimal mechanism

(b) One-sided: Prior-free and zero-reserve performance relative to the revenue optimal mechanism for a designer with marginal cost 0 capacity 5, and n buyers

Figure 4: Panel (a): Ratio of the expected weighted objective in the prior-free optimal mechanism defined in (19)–(23) to the α-optimal mechanism for various weights α on revenue. Panel (b): Ratio of the expected revenue in the one-sided prior-free optimal mechanism with $\alpha^B = 1$ and $\alpha^S = 0$ to the expected revenue in the optimal mechanism. Also shown is the ratio of the expected revenue in the one-sided $m+1$-st price procurement with no reserve to the expected revenue in the optimal mechanism. Both comparisons assume a designer with marginal cost zero and capacity of 5 units. Panels (a) and (b) are both based on Monte Carlo simulation (5000 auctions) with values and costs drawn from the Uniform distribution on $[0, 1]$.

Our analysis encompasses the one-sided case by allowing separate values of α for buyers and sellers and setting the value to zero for the side of the market with no private information. For example, consider a one-sided setup with n potential buyers and a designer with constant marginal cost of zero and capacity of m units.[48] As shown in the example of Figure 4(b), the prior-free optimal one-sided mechanism performs better than a "$m+1$-st" price auction for sufficiently small n, but provides little or no advantage when n is large. In a one-sided

[48] If we drop the assumption $\underline{v} < \overline{c}$, this fits our setup by assuming that there are m sellers without private information whose costs are "drawn" from the degenerate distribution G with support $[\underline{c}, \overline{c}]$ with $\underline{c} = 0 = \overline{c}$.

setup, knowledge of distributions is necessary to determine the optimal reserve price, but this knowledge becomes less important as the number of buyers increases because the optimal reserve price is less likely to be binding. Thus, in a one-sided setting, inferences about the underlying distribution from bid data may be of limited value when there are many bidders.