Asymptotically Optimal Prior-Free Clock Auctions*

Simon Loertscher† Leslie M. Marx‡

January 16, 2019

Abstract

Clock auctions have a number of properties that make them attractive for practical purposes. They are weakly group strategy-proof, make bidding truthfully an obviously dominant strategy, and preserve trading agents' privacy. However, optimal reserve prices and stopping rules depend on the details of underlying distributions, and so clock auctions have proved challenging to implement in a prior-free, asymptotically optimal way. In this paper, we develop a prior-free clock auction that is asymptotically optimal by exploiting 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, 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, Patrick Rey, Larry Samuelson, Mark Satterthwaite, Ilya Segal, Thomas Tröger, Rakesh Vohra, Steve Williams, Robert Wolpert, and seminar participants at the Australian National University, Northwestern University, Toulouse School of Economics, 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. We gratefully acknowledge support from the European Research Council under the Grant Agreement no. 340903, from the Samuel and June Hordern Endowment, and from the Faculty of Business and Economics at the University of Melbourne via an Eminent Research Scholar Grant.

†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 popular robustness requirements.

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.[3]

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 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 mechanism.[4]

[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 (forth.).

[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] 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, such as those by Delacrétaz et al. (2019), arise because the designer faces the constraint that revenue must not be negative.

[4] 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 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 of types drawn from the same distribution and the spacings (distances) between them and the fact that, under the regularity assumption that virtual types are increasing, no knowledge of the inframarginal virtual types is required to determine the Bayesian optimal allocation.[5] Assuming agents play their obviously dominant strategies and, on each side of the market, draw their types from identical distributions, the spacings between nontrading agents' types are given by the spacings between their exit prices and are, thus, observable. As we show, these observations and assumptions are enough for a uniformly consistent estimate of the virtual types of the marginal active buyer and seller. 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. Because no information about inframarginal types is required and because estimates are uniformly consistent, a clock auction that stops at the earliest point at which the estimated virtual type of the marginal active buyer exceeds the estimated virtual cost of the marginal active seller is asymptotically optimal.

To the best of our knowledge, ours is the first paper to develop an asymptotically optimal clock auction for a general setting in which, if distributions were known, the optimal mechanisms in the tradition of Myerson (1981) would be well understood. Furthermore, the structure of the clock auction is essentially pinned down if one requires it to be sequentially consistent in that, similar to Akbarpour and Li's (2018) notion of credible mechanisms, there is no commitment problem for the auctioneer in the dynamic implementation. 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 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 (forth.), and Li (2017), this literature has identified advantages of dynamic implementation over direct mechanisms in a variety of

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.

[5]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.

setups.[6] In particular, our paper builds on the properties of clock auctions identified by Milgrom and Segal (forth.) and on design features first introduced by McAfee (1992).[7]

Motivated by Wilson (1987), we develop prior-free clock auctions that are asymptotically optimal in the sense that, in the large, they implement the optimal mechanism derived by Myerson (1981).[8] We do so 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 (2018) 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 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

[6]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.

[7]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.

[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 (2018) 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).

[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).

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 introduces the setup, Bayesian optimality, and clock auctions. Among other things, it decribes the Bayesian optimal mechanism and shows that in a two-sided setup, the Bayesian optimal mechanism is not clock implementable. Section 3 shows by construction that a prior-free clock auction exists that is asymptotically optimal. In addition, we show that the structure of the prior-free clock auction can be pinned down by a notion of sequential consistency, and we provide criteria for determining the details of the required estimators. Section 4 discusses extensions, and Section 5 concludes.

2 Setup

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

[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.

market are privately informed about their types. In other words, we study 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 F and that all sellers draw their types from the same distribution G. In the online appendix we extend our results to allow for heterogeneity across groups of buyers and heterogeneity across groups of sellers.[14]

When considering prior-free mechanisms, we assume that each agent with private information is privately informed about his type, but the types and distributions that they are drawn from 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.[15] 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 spell out below. We also 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.[16] Individual rationality is most naturally required to be satisfied ex post.[17] These are therefore the notions that we focus on going forward.

[14] Specifically, in the extension we assume that all buyers within a group draw their values from the same distribution and that all sellers within a group draw their costs from the same distribution, but we allow the distributions to vary across groups. This extension allows for price discrimination across groups as well as the implementation of revenue thresholds, group-specific caps, and group-specific favoritism.

[15] 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.

[16] 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.

[17] 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).

2.1 Benchmark Bayesian optimal mechanisms

We now describe the Bayesian optimal mechanisms for the informational setups that we consider as well as regularity assumptions.

Bayesian optimality with one-sided private information pertaining to buyers

For the setup with *one-sided private information pertaining to buyers*, $c_{[1]}, ..., c_{[m]}$ define 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.[18] 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

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

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 $\Phi(v_{(q)}) \geq c_{[q]}$.[19] The q buyers with the highest values trade and, in the dominant strategy implementation, pay the price $p^B = \max\{v_{(q+1)}, \Phi^{-1}(c_{[q]})\}$.[20] 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 $v_{(1)}, ..., v_{(n)}$ are commonly known, $\overline{c} \geq v_{(1)} > \underline{c}$, and the virtual cost function

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

[18] 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.

[19] 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 $E_{\mathbf{v}|F,...,F}[\sum_{i=1}^{n} \Phi(v_i) q_i(\mathbf{v}, \mathbf{c}) - \sum_{j=1}^{m} c_j q_j(\mathbf{v}, \mathbf{c})]$, 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 $\sum_{i=1}^{n} q_i(\mathbf{v}, \mathbf{c}) \leq \sum_{j=1}^{m} q_j(\mathbf{v}, \mathbf{c})$.

[20] In case multiple buyers have the same value $v_{(q)}$, that is, when $v_{(q)} = v_{(q+i)}$ for $i = 1, \ldots,$ the mechanism needs to ration the buyers with this value. Any arbitrary tie-breaking rule will do without distorting incentives as the agent who wins the tie-break gets a payoff of 0 just like the agents who lose the tie-break.

is increasing, then this is a standard procurement auction in which the optimal quantity traded is the largest index q such that $v_{(q)} \geq \Gamma(c_{[q]})$.[21] The q sellers with the lowest costs trade and, in the dominant strategy implementation, are paid $p^S = \min\{c_{[q+1]}, \Gamma^{-1}(v_{(q)})\}$.[22]

Bayesian optimality with two-sided private information

For *two-sided private information*, we let the set of (privately informed) agents be $\mathcal{N} \cup \mathcal{M}$, where $\mathcal{N}$ is the set of buyers with unit demands, whose cardinality is n, and $\mathcal{M}$ 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

$$\overline{c} \geq \overline{v} > \underline{c} \geq \underline{v}. \tag{1}$$

We refer to condition (1) 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. The q buyers with the highest values and the q sellers with the lowest costs 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 $p^S = \min\{c_{[q+1]}, \Gamma^{-1}(\Phi(v_{(q)}))\}$, where ties among marginal traders can be broken arbitrarily.

Asymptotic optimality

Now that we have identified the Bayesian optimal outcome, we can define what we mean for a mechanism to be asymptotically optimal. In the tradition of Gresik and Satterthwaite (1989), we derive asymptotic results assuming independent private values by considering η-fold replicas of the economy, which ensures that the numbers of buyers and sellers go to

[21] 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 $E_{\mathbf{c}|G,\ldots,G}[\sum_{i=1}^{n} v_i q_i(\mathbf{v}, \mathbf{c}) - \sum_{j=1}^{m} \Gamma(c_j) q_j(\mathbf{v}, \mathbf{c})]$.

[22] Like with buyers, in case multiple sellers have the same cost $c_{[q]}$, an arbitrary tie-breaking rule can be applied without distorting incentives.

[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 $E_{\mathbf{v},\mathbf{c}|F,\ldots,F,G,\ldots,G}[\sum_{i=1}^{n} \Phi(v_i) q_i(\mathbf{v}, \mathbf{c}) - \sum_{j=1}^{m} \Gamma(c_j) q_j(\mathbf{v}, \mathbf{c})]$.

infinity in fixed proportion. Given an initial set of agents, an η-fold replica has a set of agents equal to the union of η instances of the initial set of agents, with the corresponding type space for each replica, and beliefs defined over a single replica of the type space. We say that an incentive compatible, individually rational mechanism defined for an η-fold replica is *asymptotically optimal* if the ratio of the designer's objective in the mechanism to the designer's objective in the Bayesian optimal mechanism converges to 1 in probability as the number of replicas goes to infinity.

Regularity assumptions

The Bayesian optimal mechanism described above requires the assumption of increasing virtual type functions, which is what Myerson (1981) refers to as the *regular case*. This assumption ensures that point-by-point maximization permits incentive compatible implementation because it implies that more efficient types—buyers with higher values and sellers with lower costs—are more likely to trade. If, for a two-sided setting, Φ is not increasing in the neighborhood of some v' with $\Phi(v') > \underline{c}$ and/or Γ is not increasing 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.

For our purposes, we can relax the usual regularity condition considerably. To illustrate, consider 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 among all prices $(\hat{p}^B, \hat{p}^S) \in [\underline{v}, \overline{v}] \times [\underline{c}, \overline{c}]$ in the limit as we replicate an initial set of n buyers and m sellers, i.e., "in the large," is that the prices satisfy

$$\Phi(p^B) = \Gamma(p^S) \quad \text{and} \quad n(1 - F(p^B)) = mG(p^S). \tag{2}$$

If the solution to (2) is unique and if the Bayesian optimal mechanism in the large is a posted-price mechanism (both of which are, for example, the case when the virtual type functions are increasing), then the mechanisms described above are Bayesian optimal in the large even when they are not Bayesian optimal in the small (see Figure 1 for an example).[25]

Correspondingly, our results for asymptotically optimal prior-free mechanisms require only the assumption that there is a unique pair of prices satisfying (2) and that the Bayesian optimal mechanism in the large is a posted-price mechanism. In what follows, we maintain this assumption.

The tight connection between the Bayesian optimality of posted-price mechanisms in the

[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 (2) by replacing Γ and Φ by the identity function.

large and the asymptotic optimality of our prior-free 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 (2) 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 (2) holds, with Φ and Γ replaced by prior-free estimates.

As an illustration, Figure 1 depicts a case in which, in the large with equal numbers of buyers and sellers, the Bayesian optimum is implemented with posted prices p^B and p^S such that the 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. Further, if the example 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 posted-price mechanism because some share of agents would have to be rationed.

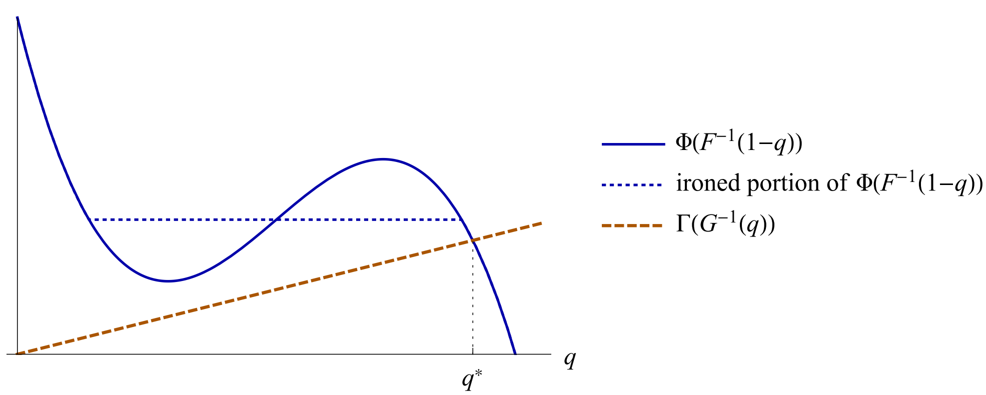

Figure 1: Illustration of a setup with unique posted prices that are Bayesian optimal in the large, but where ironing may be required in the small.

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] with weight $\alpha \in [0, 1]$ on expected profit, 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*

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

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. However, to allow for the possibility that the virtual type functions are not increasing, and so not invertible, we define the inverse virtual value function to select the lowest corresponding value and the inverse virtual cost function to select the largest corresponding cost. Formally,

$$\Phi_\alpha^{-1}(x) \equiv \min\{v \mid \Phi_\alpha(v) \geq x\} \quad \text{and} \quad \Gamma_\alpha^{-1}(x) \equiv \max\{c \mid \Gamma_\alpha(c) \leq x\}.$$

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.

2.2 Clock auctions

Next, we define clock auctions and discuss some of their key characteristics.

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

[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)). Alternatively, and equivalently, if λ is the solution value of the Lagrange-mulitplier associated with the problem of maximizing expected social surplus subject to the constraint of generating some minimum expected profit (and subject to incentive compatibility and individual rationality constraints), the optimal allocation rule is as just described with $\alpha = \lambda/(1+\lambda)$.

in Milgrom and Segal (forth.) to accommodate (without requiring) a two-sided setting.[28] In 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 $\phi : \Omega \to \mathbb{R}$, which is increasing in the buyer clock price, a seller function $\gamma : \Omega \to \mathbb{R}$, which is increasing in the seller clock price, and a target function $\tau : \Omega \to \mathbb{R}$, 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, 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.

[28] Milgrom and Segal (forth.) 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 the active set.

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

Key properties of clock auctions

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

In our environment with single-unit demands and single-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.

Clock auctions are also *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.[29] In addition, a clock auction is *envy free* in the sense that in equilibrium no agent prefers the ex post allocation and price paid by or to another agent to his own.[30]

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

[29]On weak group strategy-proofness in a one-sided clock auction, see Li (2017) and Milgrom and Segal (forth.). 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.

[30]This property is not unique to the clock auction implementation. The dominant strategy implementations of the Bayesian optimal mechanisms derived above are also envy free.

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.

Achieving the Bayesian optimum in a 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 *In the setup with one-sided private information, the Bayesian optimal mechanism can be implemented by a clock auction if the virtual type function for the side with private information is increasing. In the setup with two-sided private information, the Bayesian optimal mechanism cannot be implemented by a clock auction with finite numbers of agents.*

Proposition 1 summarizes the implication of the results of Milgrom and Segal (forth.) for our setting.[31] 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.[32] 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.

3 Asymptotically optimal clock auctions

We now show the existence of a prior-free clock auction that is asymptotically optimal. To have a clock auction, we require a buyer function, seller function, and target function, which can only be estimated based on the types of agents that have exited and so do not trade. The requirement of asymptotic optimality places constraints on how these functions are estimated. In addition, we show that the structure of the prior-free clock auction can

[31]For settings like ours, Proposition 6 in Milgrom and Segal (forth.) 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 for each other, whereas when private information pertains to both sides, buyers and sellers are complements—the problem is assignment representable, as defined by Delacrétaz et al. (2019), and so the complementarity of buyers and sellers follows from Shapley (1962).

[32]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 $\tau(\boldsymbol{\omega}) \equiv c_{[n^A]}$.

be 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 below.

3.1 Consistent estimation in clock auctions

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 1 *For all $j \in \{1, ..., n-1\}$,*

$$jE_{\mathbf{v}|F,...,F}[v_{(j)} - v_{(j+1)}] = E_{\mathbf{v}|F,...,F}\left[\frac{1-F(v_{(j)})}{f(v_{(j)})}\right]$$

and for all $j \in \{1, ..., m-1\}$,

$$jE_{\mathbf{c}|G,...,G}[c_{[j+1]} - c_{[j]}] = E_{\mathbf{c}|G,...,G}\left[\frac{G(c_{[j]})}{g(c_{[j]})}\right].$$

Proof. See Appendix B.

Lemma 1 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 asymptotic optimality of the prior-free clock auction defined below. We ignore the target function for the moment because it only affects the number of trades by at most one and so does not affect the asymptotic properties of the mechanism. 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 define the buyer and seller functions as follows:

$$\phi(\boldsymbol{\omega}) = p^B - \alpha j \sigma_j^v \text{ and } \gamma(\boldsymbol{\omega}) = p^S + \alpha j \sigma_j^c, \tag{3}$$

where σ_j^v and σ_j^c are spacing estimators. Specifically, 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. 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 σ_j^v be the distance between $\hat{v}_{(j)}$ and the r_n-th nearest neighbor less than σ_j^v, divided by r_n, and similarly for σ_j^c. More precisely, to account for the possibility that r_n exceeds the number of available observations, we define the spacing estimators as follows:

$$\sigma_j^v \equiv \begin{cases} \frac{\hat{v}_{(j)} - \hat{v}_{(j+\min\{r_n, n-j\})}}{\min\{r_n, n-j\}}, & \text{if } j < n \\ \frac{1}{n+1}, & \text{otherwise} \end{cases} \quad \text{and } \sigma_j^c \equiv \begin{cases} \frac{\hat{c}_{[j+\min\{r_m, m-j\}]} - \hat{c}_{[j]}}{\min\{r_m, m-j\}}, & \text{if } j < m \\ \frac{1}{m+1}, & \text{otherwise,} \end{cases} \tag{4}$$

where r_n (and correspondingly r_m) satisfies

$$\lim_{n\to\infty} r_n = \infty \quad \text{and} \quad \lim_{n\to\infty} \frac{r_n}{n} = 0 \quad \text{and} \quad \lim_{n\to\infty} \frac{\ln n}{r_n} = 0. \tag{5}$$

For example, the conditions in (5) are satisfied if $r_n = n^x$ for $x \in (0,1)$.

As we now show, $j\sigma_j^v$ provides a uniformly consistent estimator for $\frac{1-F(v_{(j)})}{f(v_{(j)})}$ and $j\sigma_j^c$ provides a uniformly consistent estimator for $\frac{G(c_{[j]})}{g(c_{[j]})}$. Our use of $j\sigma_j^v$ as an estimate of $\frac{1-F(v_{(j)})}{f(v_{(j)})}$ means that if we take the empirical cumulative distribution as the estimate of $F(v_{(j)})$, then we have $j\sigma_j^v = \frac{j/n}{\hat{f}(v_{(j)})}$, where $\hat{f}$ is our implied density estimator and thus

$$\hat{f}(v_{(j)}) = \frac{\frac{j}{n}}{j\sigma_j^v} = \frac{1}{n\sigma_j^v} = \frac{r_n}{n(v_{(j)} - v_{(j+r_n)})}.$$

It follows that $\hat{f}(v_{(j)})$ is a one-sided Loftsgaarden-Quesenberry (1965) nearest neighbor estimator for the density f when evaluated at $v_{(j)}$.

By the Glivenko-Cantelli theorem, the empirical cumulative distribution converges uniformly in probability to the actual cdf, and so the uniform convergence in probability of our density estimator, as shown in Theorem 1, completes the proof. Although Theorem 1 has, to our knowledge, not been shown previously, and its proof is fairly complex, the result is at its heart a straightforward adaptation of the result of Moore and Henrichon (1969). The theorem of Moore and Henrichon (1969), which draws on Loftsgaarden and Quesenberry (1965), shows the uniform consistency of a two-sided nearest neighbor estimator. We adapt the result and associated proof to the case of a one-sided nearest neighbor estimator that is evaluated only at observed values.

Theorem 1 *Given that F is continuously differentiable with positive density f on support*

$[\underline{v}, \overline{v}]$ and that $r_n \to \infty$, $\frac{r_n}{n} \to 0$, and $\frac{\ln n}{r_n} \to 0$, it follows that as $n \to \infty$,

$$\Pr\left[\sup_{\ell \in \{1,\ldots,n-r_n\}} \left| \ell \sigma_\ell^v - \frac{1 - F(v_{(\ell)})}{f(v_{(\ell)})} \right| > \varepsilon \right] \to 0.$$

Analogously, given that G is continuously differentiable with positive density g on support $[\underline{c}, \overline{c}]$, it follows that as $m \to \infty$,

$$\Pr\left[\sup_{\ell \in \{1,\ldots,m-r_m\}} \left| \ell \sigma_\ell^c - \frac{G(c_{[\ell]})}{g(c_{[\ell]})} \right| > \varepsilon \right] \to 0.$$

Proof. See Appendix B.

To see the intuition for Theorem 1, note that by the definition of a density, our differentiability assumptions, and the result that $r_n \sigma_\ell^v$ goes to zero with n, we have

$$\Pr\left(\sup_{\ell \in \{1,\ldots,n-r_n\}} \left| \frac{(F(v_{(\ell)}) - F(v_{(\ell)} - r_n \sigma_\ell^v))}{r_n \sigma_\ell^v} - f(v_{(\ell)}) \right| > \varepsilon \right) \to 0.$$

Also, as shown in the proof of Theorem 1,

$$\Pr\left(\sup_{\ell \in \{1,\ldots,n-r_n\}} \left| \frac{n(F(v_{(\ell)}) - F(v_{(\ell)} - r_n \sigma_\ell^v))}{r_n} - 1 \right| > \varepsilon \right) \to 0,$$

which holds because $F(v_{(\ell)}) - F(v_{(\ell)} - r_n \sigma_\ell^v)$ has the same asymptotic properties as the sum of r_n out of $n+1$ exponential random variables with mean 1 divided by the sum of all $n+1$ exponential random variables. Loosely speaking, the sum of the r_n exponential random variables with mean 1 in the numerator "cancels" with the r_n in the denominator, and the n in the numerator "cancels" with the sum of $n+1$ exponential random variables with mean 1 in the denominator. Thus, the ratio stays close to 1 for n sufficiently large. Putting these together, we have

$$\Pr\left(\sup_{\ell \in \{1,\ldots,n-r_n\}} \left| \frac{r_n}{n} \frac{1}{r_n \sigma_\ell^v} - f(v_{(\ell)}) \right| > \varepsilon \right) \to 0,$$

which completes the proof because $\frac{r_n}{n} \frac{1}{r_n \sigma_\ell^v} = \frac{1}{n \sigma_\ell^v} = \hat{f}(v_{(\ell)})$, which is our implied density estimator.

Using Theorem 1, it follows that $\phi(v_{(j)})$ and $\gamma(c_{[j]})$ as defined in (3) are uniformly consistent estimators of $\Phi_\alpha(v_{(j-1)})$ and $\Gamma_\alpha(c_{[j-1]})$, respectively, away from the boundary, i.e., for a number of replicas η of an initial set of n buyers and m sellers such that $r_{\eta n} < \eta n - j$ and $r_{\eta m} < \eta m - j$. Because $\Phi_\alpha(\underline{v}) \leq \underline{v} < \overline{c} \leq \Gamma_\alpha(\overline{c})$, it follows that as the number of replicas of the economy goes to infinity, the Bayesian optimal number of trades is away from the boundary in the sense of being less than $\min\{\eta n - r_{\eta n}, \eta m - r_{\eta m}\}$ with probability 1. Thus, we have the following result:

Proposition 2 *The prior-free clock auction defined by* (3)–(5)*, combined with any target function, is asymptotically optimal.*

The adaptation of Proposition 2 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 $c_{[j]}$.

The foundation for Proposition 2 is the result that as the number of agents goes to infinity, the estimated virtual types defined by (3)–(5) are close to the theoretical virtual types, so the "first time" (as agents exit a clock auction) that the estimated virtual types cross, with the estimated virtual value exceeding the estimated virtual cost, cannot be too far from the one and only time that the theoretical virtual types cross. Our estimated virtual type functions differ from the "empirical" virtual type functions in that we use a uniformly convergent (in probability) estimator for the density based on the average of number of adjacent spacings between worse types rather than a local estimate of the rate of change in the empirical distribution. Given that the local empirical estimate of $g(c_{[j]})$ is $\frac{\frac{j+1}{m}-\frac{j}{m}}{c_{[j+1]}-c_{[j]}}$, the local empirical virtual cost is $\hat{\Gamma}(c) \equiv c + \frac{\frac{j}{m}(c_{[j+1]}-c_{[j]})}{\frac{j+1}{m}-\frac{j}{m}} = c + j(c_{[j+1]} - c_{[j]})$, and, analogously, the local empirical virtual value is $\hat{\Phi}(v) \equiv v - j(v_{(j)} - v_{(j+1)})$. As Figure 2 illustrates, the high volatility of these local empirical virtual types renders them unhelpful for developing an asymptotically optimal prior-free clock auction; however, a clock auction based on estimates derived from the average of an appropriately chosen number of nearby spacings is asymptotically optimal.

Figure 2 considers the case of revenue maximization ($\alpha = 1$) and $n = 1000$ agents whose types are drawn from the uniform distribution on $[0, 1]$. As shown in panel (a), the theoretical virtual types indicate an optimal quantity of 250 trades. However, the first intersection of the local empirical virtual types (the first intersection of the upper and lower "clouds" in the figure) occurs at approximately twice as large a quantity. Indeed, the volatility of the those estimates pulls the implied quantity traded towards the efficient quantity of 500 rather than the optimal quantity of 250. In contrast, as shown in panel (b), the estimated virtual types based on (3)–(5) provide a much better approximation of the theoretical virtual types in that example and deliver an approximately optimal number of trades.

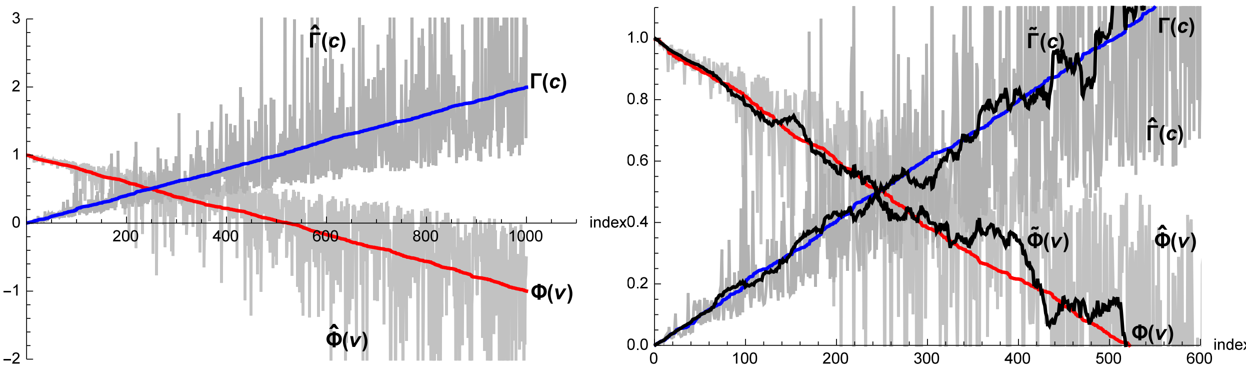

(a) Empirical virtual types

(b) Estimated virtual types based on (3)–(5)

Figure 2: Panel (a): Theoretical virtual types Φ and Γ and local empirical virtual types $\hat{\Phi}$ and $\hat{\Gamma}$. Panel (b): Theoretical virtual types Φ and Γ, local empirical virtual types $\hat{\Phi}$ and $\hat{\Gamma}$, and estimated virtual types $\tilde{\Phi}$ and $\tilde{\Gamma}$ based on (3)–(5). Both panels are based on 1000 values and 1000 costs drawn from the uniform distribution on $[0, 1]$.

3.2 Sequential consistency

In this subsection, we discuss the implication of sequential consistency for the virtual type estimator.

As shown in Proposition 2, the prior-free clock auction defined in (3)–(5) is asymptotically optimal. However, the asymptotic optimality result would continue to hold for any specification of ϕ and γ of the form:

$$\phi(\boldsymbol{\omega}) = p^B - \chi_{\alpha,j}\sigma_j^v \text{ and } \gamma(\boldsymbol{\omega}) = p^S + \chi_{\alpha,j}\sigma_j^c, \tag{6}$$

where $\chi_{\alpha,j}$ is nonnegative (ensuring that the mechanism is deficit free) and satisfies

$$\lim_{j\to\infty} \frac{\chi_{\alpha,j}}{j} = \alpha,$$

so that $\chi_{\alpha,j}$ has the asymptotic properties of αj. This raises the question whether $\chi_{\alpha,j}$ can be pinned down if one imposes additional restrictions.

We show that the answer is affirmative if one requires a prior-free analogue of sequential rationality that we call *sequential consistency*. In particular, we show that the sequential consistency criterion defined below is satisfied if and only if, for $j \in \{1, ..., \min\{m, n\}\}$,

$$\chi_{\alpha,j} \equiv \max\{0,\ \alpha(j-2) - (1-\alpha)\}. \tag{7}$$

To define sequential consistency, assume that the designer delegates the operation of a dynamic mechanism to a decision maker, such as an auctioneer. We consider whether it is credible that an expected-payoff maximizing auctioneer would follow the protocol defined by the mechanism. This is similar to the credibility notion developed independently by Akbarpour and Li (2018). Sequential consistency differs from their definition of a credible extensive form game plus strategy profile by specifying how the auctioneer forms expectations. Our notion 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.[33]

Consider 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 $\mathbf{v}_{(j)} \equiv (v_{(j)}, ..., v_{(n)})$ and $\mathbf{c}_{[j]} \equiv (c_{[j]}, ..., c_{[m]})$. 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 the beliefs generated by σ_j^v and σ_j^c that the next value and cost will be $v_{(j-1)} = v_{(j)} + \sigma_j^v$ 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 say that a prior-free clock auction is *sequentially consistent* with respect to spacing estimates $\boldsymbol{\sigma}^v$ and $\boldsymbol{\sigma}^c$ if (i) and (ii) hold.

The prior-free clock auction with ϕ and γ defined in (6) satisfies this condition if

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

which using Lemma 1, 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

$$v_{(j)} - (\alpha(j-2) - (1-\alpha))\sigma_j^v \geq c_{[j]} + (\alpha(j-2) - (1-\alpha))\sigma_j^c. \tag{8}$$

[33]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.

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 (8) hold if and only if

$$p^B - \chi_{\alpha,j}\sigma_j^v \geq p^S + \chi_{\alpha,j}\sigma_j^c, \tag{9}$$

where $\chi_{\alpha,j}$ is defined in (7). Because (9) is the criterion for ending the auction in the prior-free optimal clock auction defined by (6)–(7), the auctioneer's incentives are aligned with the auction protocol.

Proposition 3 *The stopping rule for the auctioneer in the prior-free clock auction defined in* (6)–(7) *is sequentially consistent with respect to 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.*

3.3 Target functions

To complete the definition of our asymptotically optimal prior-free clock auction, it remains to specify a target function. We define a target estimator as follows: 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., that 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 an assumption on the distributions of $v_{(j-1)}$ and $c_{[j-1]}$.

For example, the "uniform target estimator" assumes that $v_{(j-1)}$ is distributed uniformly between $v_{(j)}$ and $v_{(j)} + 2\sigma_j^v$ and that $c_{[j-1]}$ is distributed uniformly between $c_{[j]} - 2\sigma_j^c$ and $c_{[j]}$. Thus, the uniform target estimator is given by:[34]

$$\tau(\boldsymbol{\omega}) = \min\left\{\gamma(\boldsymbol{\omega}),\ \max\left\{\phi(\boldsymbol{\omega}),\ \frac{\phi(\boldsymbol{\omega}) + \gamma(\boldsymbol{\omega})}{2} + \left(1 - \frac{\alpha}{2}\right)\left(\sigma_j^v - \sigma_j^c\right)\right\}\right\}. \tag{10}$$

According to (10), 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

[34] 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 $c_{[j]}$.

above, McAfee (1992) uses the midpoint between the standing clock prices as the target, which corresponds to the midpoint between $v_{(j)}$ and $c_{[j]}$.

3.4 Discussion

In this subsection, we discuss the calculation of confidence intervals for the estimated gain from continuing an auction, criteria for selecting the number of observations to be included in the spacing estimator, and the interaction between clock auctions and the estimation of virtual type functions.

Confidence intervals

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).[35] This is illustrated in Figure 3 for an example with 20 buyers and 20 sellers.[36] Panel (a) shows 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.

As can be seen from Figure 3(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 clock auction (setting aside the target function for purposes of the illustration), the auction would end

[35] 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.

[36] 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 empirical cdf and nearest neighbor density estimator. 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.

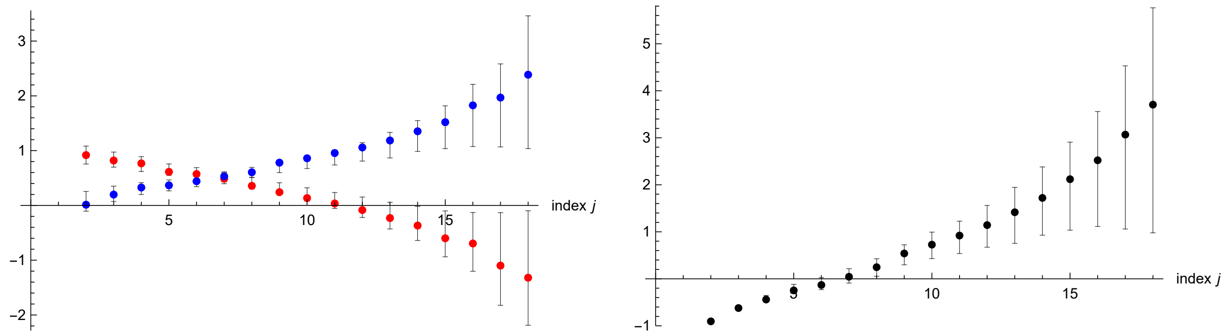

(a) Estimated virtual types with CIs

(b) Estimated revenue changes with CIs

Figure 3: 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].

following this exit. Turning to Figure 3(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 spacing estimator

As a matter of statistics, there is some flexibility in the precise details of the spacing estimator used in an asymptotically optimal prior-free clock auction. The nearest neighbor estimators that we use have bias and variance that go to zero as the initial set of agents is replicated when r_n satisfies (5), which, as noted above, holds if $r_n = n^x$ for $x \in (0, 1)$.

Further, nearest neighbor estimators have a mean square error of order $(r_n/n)^4 + 1/r_n$, which is minimized when r_n is proportional to $n^{4/5}$ (Silverman, 1986, Chapters 3 and 5.2.2), which suggests the use of $r_n = n^{4/5}$. In this case, the approximate value of the mean integrated square error tends to zero at the rate $n^{-4/5}$ (Silverman, 1986, Chapter 3.7.2), which we illustrate in the online appendix.

It follows that the nearest neighbor estimators that attain the minimum mean square error are uniquely defined up to proportionality constants.

Clock auctions and estimation

As shown in Proposition 1, clock auctions are not without loss of generality in the two-sided setup. This is because the allocation and payments in the Bayesian optimal mechanism depend on the types of the marginal trading agents, i.e., on the value of the lowest valuing buyer who trades and the cost of the highest cost seller who trades. Because these types belong to trading agents, information about them is not available in a clock auction. Interestingly, as we now show, for consistent estimation of virtual types (which, as mentioned, is necessary for asymptotic optimality), clock auctions are without loss of generality for asymptotically optimal, dominant strategy mechanisms that are envy free in a sense that we are now going to make precise.

To that end, consider first the Bayesian optimal mechanism when the distributions, and hence the virtual type functions, are known. With known distributions, there is, obviously, no need to estimate the virtual type functions, and the mechanism needs only elicit the information necessary to evaluate virtual types in order to induce the most efficient traders, according to virtual types, to trade. In contrast, a prior-free asymptotically optimal mechanism needs to first *estimate* the virtual type functions before evaluating these functions in order to rank the realized virtual types.[37] As we now argue, the data that can be used for estimation in an asymptotically optimal clock auction are the same as the data that can be used in any asymptotically optimal mechanism that satisfies dominant strategy incentive compatibility and envy-freeness. It is in this sense that clock auctions are without loss of generality with regard to estimation.[38]

Suppose that we have an asymptotically optimal prior-free direct mechanism that is dominant strategy inceyentive compatible and envy free. Suppose further that for a given realization of types, the mechanism induces trade by k buyers and k sellers. By envy-freeness and dominant strategy incentive compatibility, the traders are the k highest valuing buyers and the k lowest cost sellers, with each trading buyer paying the same threshold type (that is, the lowest value such that he still trades), and each trading seller being paid the same threshold type (that is, the highest cost such that she still trades).

From dominant strategy incentive compatibility follows further that the buyers' common threshold type cannot depend on the type of any trading buyer, and the sellers' common threshold type cannot depend on the type of any trading seller. Thus, there must be a threshold for the buyers $T_k^B(v_{(k+1)}, ..., v_{(n)}; \mathbf{c})$ such that each of the k trading buyers pays

[37] There is an apparent arbitrariness when saying that a prior-free asymptotically optimal mechanism needs to both estimate and evaluate virtual type functions because one could also view estimation and evaluation as a single step. However, the terminology makes precise the parallels and differences to Bayesian mechanisms, where there is no need to estimate virtual type functions.

[38] The restriction to envy-free mechanisms deserves a brief elaboration. Although the dominant strategy implementation of the Bayesian optimal mechanism satisfies envy-freeness, there are asymptotically optimal mechanisms—such as the one of Baliga and Vohra (2003)—that are not envy free away from the limit. Consequently, these mechanisms can use data for estimation that are not available in a clock auction.

$T_k^B(v_{(k+1)},...,v_{(n)};\mathbf{c})$ and such that each trading buyer would no longer trade if, holding other types fixed, he had a value less than $T_k^B(v_{(k+1)},...,v_{(n)};\mathbf{c})$. Likewise, there must be a threshold for the sellers $T_k^S(c_{[k+1]},...,c_{[m]};\mathbf{v})$ such that each of the k trading sellers is paid $T_k^S(c_{[k+1]},...,c_{[m]};\mathbf{v})$ and such that each trading seller would no longer trade if, holding other types fixed, she had a cost greater than $T_k^S(c_{[k+1]},...,c_{[m]};\mathbf{v})$.

Asymptotic optimality of a prior-free mechanism implies that the data available to the mechanism must be used to estimate the virtual value and cost functions. Specifically, assume that the quantity traded is k and let $\hat{\Phi}_k(\cdot \mid \mathbf{x}_k^B)$ and $\hat{\Gamma}_k(\cdot \mid \mathbf{x}_k^S)$ be the estimated virtual type functions that are used to evaluate, respectively, $v_{(k)}$ and $c_{[k]}$, where $\mathbf{x}_k^B$ is the data used to estimate the function $\hat{\Phi}_k$ and $\mathbf{x}_k^S$ is the data used to estimate the function $\hat{\Gamma}_k$. By incentive compatibility, these functions are increasing in $v_{(k)}$ and $c_{[k]}$, respectively, and hence invertible. It follows that the threshold types are $\max\{v_{(k+1)},\hat{\Phi}_k^{-1}(\hat{\Gamma}_k(c_{[k]};\mathbf{x}_k^S);\mathbf{x}_k^B)\}$ for trading buyers and $\min\{c_{[k+1]},\hat{\Gamma}_k^{-1}(\hat{\Phi}_k(v_{(k)};\mathbf{x}_k^B);\mathbf{x}_k^S)\}$ for trading sellers. But now we have

$$\max\{v_{(k+1)},\hat{\Phi}_k^{-1}(\hat{\Gamma}_k(c_{[k]};\mathbf{x}_k^S);\mathbf{x}_k^B)\} = T_k^B(v_{(k+1)},...,v_{(n)};\mathbf{c}),$$

where $T_k^B(v_{(k+1)},...,v_{(n)};\mathbf{c})$ does not depend on $v_{(1)},...,v_{(k)}$, and

$$\min\{c_{[k+1]},\hat{\Gamma}_k^{-1}(\hat{\Phi}_k(v_{(k)};\mathbf{x}_k^B);\mathbf{x}_k^S)\} = T_k^S(c_{[k+1]},...,c_{[m]};\mathbf{v}),$$

where $T_k^S(c_{[k+1]},...,c_{[m]};\mathbf{v})$ is independent of $c_{[1]},...,c_{[k]}$. It follows that $\mathbf{x}_k^S$ and $\mathbf{x}_k^B$ do not include $v_{(1)},...,v_{(k)}$ and $c_{[1]},...,c_{[k]}$. In other words, the data $\mathbf{x}_k^B$ and $\mathbf{x}_k^S$ used to estimate $\hat{\Phi}_k$ and $\hat{\Gamma}_k$ include only the types of nontrading agents.

Let us say that a mechanism is *defined by virtual type functions* if, given (possibly estimated) virtual type functions, it trades the largest number k such that the virtual value associated with $v_{(k)}$ is greater than or equal to the virtual cost associated with $c_{[k]}$. Further, we say that two mechanisms are *asymptotically equivalent* if, in the limit as the number of replicas of the initial set of agents goes to infinity, the share of agents who trade and the expected payments for trading agents are the same. With this at hand, we have the following result:

Proposition 4 *Given an asymptotically optimal prior-free direct mechanism that is dominant strategy incentive compatible and envy free, there exists an asymptotically equivalent mechanism with those same properties that is defined by virtual type functions that are estimated based only on the types of nontrading agents.*

Because the types of the nontrading agents are data that are available in a clock auction, Proposition 4 implies that clock auctions are not a restriction for the estimation of virtual type functions for an asymptotically optimal prior-free mechanism that is dominant strategy incentive compatible and envy free.

4 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.

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. Similarly, in a procurement setting, firms may differ with respect to their cost distributions, for example, because some firms are the results of mergers (see e.g. Loertscher and Marx, 2019).

In the online appendix, we show how our analysis and clock auction generalize to such settings. As described there, 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 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 that group of 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. 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)). Favoritism can be accomplished in a discriminatory clock auction by assigning favored and non-favored agents to separate groups, evaluating non-favored agents using virtual types that incorporate the designer's unconstrained weight α on revenue, and evaluating favored agents using virtual types with weight α^f on revenue. Favoritism then simply means that $\alpha^f < \alpha$.

5 Conclusions

We develop a prior-free clock auction that is 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 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 to 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 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 $\boldsymbol{\omega}_t^B = (\mathcal{N}^A, \mathbf{x}^B, p^B)$ and $\boldsymbol{\omega}_t^S = (\mathcal{M}^A, \mathbf{x}^S, p^S)$ are buyer and seller states. The components of the buyer state are: the set of active buyers $\mathcal{N}^A \subseteq \mathcal{N}$ with cardinality n^A, the vector of exit prices for non-active buyers $\mathbf{x}^B \in \mathbb{R}^{n-n^A}$, and the buyer clock price $p^B \in \mathbb{R}$. 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 $\boldsymbol{\omega}_0^B = (\mathcal{N}, \emptyset, \underline{p})$ and $\boldsymbol{\omega}_0^S = (\mathcal{M}, \emptyset, \overline{p})$ 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 $\phi : \Omega \to \mathbb{R}$ is increasing in p^B, that $\gamma : \Omega \to \mathbb{R}$ is increasing in p^S, and that $\tau : \Omega \to \mathbb{R}$ 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 $\boldsymbol{\omega}_t' = (z_t, (\mathcal{N}^A, \mathbf{x}^B, T^B(\boldsymbol{\omega}_t)), \boldsymbol{\omega}_t^S)$, 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 $\boldsymbol{\omega}_t'' = (z_t, \boldsymbol{\omega}_t^B, (\mathcal{M}^A, \mathbf{x}^S, T^S(\boldsymbol{\omega}_t)))$.

For $t \in \{0, 1, ...\}$, if $\boldsymbol{\omega}_t^B = (\mathcal{N}^A, \mathbf{x}^B, p^B)$, $\boldsymbol{\omega}_t^S = (\mathcal{M}^A, \mathbf{x}^S, p^S)$, 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 $\boldsymbol{\omega}_{t+1}^B = (\mathcal{N}^A \backslash \{i\}, (\mathbf{x}^B, \hat{p}^B), \hat{p}^B)$, or the buyer clock price reaches $T^B(\boldsymbol{\omega}_t)$ with no exit, in which case $\boldsymbol{\omega}_{t+1}^B = (\mathcal{N}^A, \mathbf{x}^B, T^B(\boldsymbol{\omega}_t))$. Decrease the seller clock price from p^S until either a seller j exits at $\hat{p}^S$, in which case $\boldsymbol{\omega}_{t+1}^S = (\mathcal{M}^A \backslash \{j\}, (\mathbf{x}^S, \hat{p}^S), \hat{p}^S)$, or the seller clock price reaches $T^S(\boldsymbol{\omega}_t)$ with no exit, in which case $\boldsymbol{\omega}_{t+1}^S = (\mathcal{M}^A, \mathbf{x}^S, T^S(\boldsymbol{\omega}_t))$. 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 $\boldsymbol{\omega}_{t+1}^B = (\mathcal{N}^A \backslash \{i\}, (\mathbf{x}^B, \hat{p}^B), \hat{p}^B)$, $\boldsymbol{\omega}_{t+1}^S = \boldsymbol{\omega}_t^S$, and $z_{t+1} = 0$, or the buyer clock price reaches $\overline{p}$, in which case $\boldsymbol{\omega}_{t+1}^B = (\hat{\mathcal{N}}^A, \mathbf{x}^B, \overline{p})$, where $\hat{\mathcal{N}}^A$ consists of m^A randomly selected elements of $\mathcal{N}^A$, $\boldsymbol{\omega}_{t+1}^S = \boldsymbol{\omega}_t^S$, 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 $\boldsymbol{\omega}_{t+1}^B = \boldsymbol{\omega}_t^B$, $\boldsymbol{\omega}_{t+1}^S = (\mathcal{M}^A \backslash \{j\}, (\mathbf{x}^S, \hat{p}^S), \hat{p}^S)$, and $z_{t+1} = 0$, or the seller clock price reaches $\underline{p}$, in which case $\boldsymbol{\omega}_{t+1}^S = (\hat{\mathcal{M}}^A, \mathbf{x}^S, \underline{p})$, where $\hat{\mathcal{M}}^A$ consists of n^A randomly selected elements of $\mathcal{M}^A$, $\boldsymbol{\omega}_{t+1}^B = \boldsymbol{\omega}_t^B$, 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 $\boldsymbol{\omega}_t^B = (\mathcal{N}^A, \mathbf{x}^B, p^B)$ 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 $\boldsymbol{\omega}_{t+1}^B = (\mathcal{N}^A \backslash \{i\}, (\mathbf{x}^B, \hat{p}^B), \hat{p}^B)$ and $z_{t+1} = 0$, or the buyer clock price reaches $T^B(\boldsymbol{\omega}_t)$ with no exit, in which case $\boldsymbol{\omega}_{t+1}^B = (\mathcal{N}^A, \mathbf{x}^B, T^B(\boldsymbol{\omega}_t))$ 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 Lemma 1. 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 Theorem 1. We present the proof for the distribution of values. The proof for the distribution of costs follows by analogous arguments.

Define $d_{r_n}(z)$ to be the distance from z to the r_n-th closest of the observations among $V_1, ..., V_n$, which are iid random draws from F, that are less than or equal to z. Further, define

$$U_{r_n}(z) \equiv F(z) - F(z - d_{r_n}(z)).$$

Part 1: We show first that as $n \to \infty$,

$$\Pr\left[\sup_{\ell \in \{1,...,n-r_n\}} \left|\frac{n}{r_n} U_{r_n}(V_{(\ell)}) - 1\right| > \varepsilon\right] \to 0. \tag{11}$$

By the definition of d_{r_n}, the interval $\left[V_{(\ell)} - d_{r_n}(V_{(\ell)}), V_{(\ell)}\right]$ contains $r_n + 1$ observations, including observations at each endpoint of the interval, and r_n spacings. In particular, the lower end of the interval is $V_{(\ell+r_n)}$. Thus, for $\ell \in \{1, ..., n - r_n\}$,

$$\sum_{j=1}^{r_n} \left(F(V_{(\ell+j-1)}) - F(V_{(\ell+j)})\right) = F(V_{(\ell)}) - F(V_{(\ell+r_n)}) = F(V_{(\ell)}) - F(V_{(\ell)} - d_{r_n}(V_{(\ell)})) = U_{r_n}(V_{(\ell)}).$$

Of course, then

$$\frac{n}{r_n} \sum_{j=1}^{r_n} \left(F(V_{(\ell+j-1)}) - F(V_{(\ell+j)})\right) = \frac{n}{r_n} U_{r_n}(z).$$

It is well known (see, e.g., Nagaraja et al., 2015) that, letting $F(V_{(n+1)}) \equiv 0$ and $F(V_{(0)}) \equiv 1$, the $n + 1$ random variables,

$$F(V_{(n)}), F(V_{(n-1)}) - F(V_{(n)}), ..., F(V_{(1)}) - F(V_{(2)}), 1 - F(V_{(1)})$$

have the same joint distribution as

$$\frac{Y_{n+1}}{S_{n+1}}, ..., \frac{Y_1}{S_{n+1}},$$

where $Y_1, ..., Y_{n+1}$ are independent exponential random variables with mean 1 and $S_{n+1} = Y_1 + ... + Y_{n+1}$. This means that $\sum_{j=1}^{r_n} \left(F(V_{(\ell+j-1)}) - F(V_{(\ell+j)})\right)$ has the same distribution as $\sum_{j=1}^{r_n} \frac{Y_{\ell+j}}{S_{n+1}} = \sum_{j=\ell+1}^{\ell+r_n} \frac{Y_j}{S_{n+1}}$. Thus, $\frac{n}{r_n} U_{r_n}(z)$ will converge uniformly in probability to 1 (which would prove (11)) if we can prove that $\frac{n}{r_n} \sum_{j=\ell+1}^{\ell+r_n} \frac{Y_j}{S_{n+1}}$ converges uniformly (over all values of $\ell \in \{1, ..., n - r_n\}$) in probability to 1, i.e., that for all $\varepsilon > 0$,

$$\Pr\left(\max_{\ell\in\{1,...,n-r_n\}} \left| \frac{n}{r_n} \sum_{j=\ell+1}^{\ell+r_n} \frac{Y_j}{S_{n+1}} - 1 \right| > \varepsilon \right) \to 0, \tag{12}$$

which we can rewrite as

$$\Pr\left(\max_{\ell\in\{1,...,n-r_n\}} \left| \frac{1}{S_{n+1}/n} \sum_{j=\ell+1}^{\ell+r_n} \frac{Y_j}{r_n} - 1 \right| > \varepsilon \right) \to 0, \tag{13}$$

Because S_{n+1} is the sum of $n+1$ iid random variables with mean 1, by the law of large numbers $\frac{S_{n+1}}{n}$ converges almost surely to 1, i.e., $\Pr(\lim_{n\to\infty} \frac{S_{n+1}}{n} = 1) = 1$. Thus, (13) will follow if we can show that the sums $\sum_{j=\ell+1}^{\ell+r_n} \frac{Y_j}{r_n}$ are uniformly (for all $\ell \in \{1, ..., n - r_n\}$) near 1 in probability. For any $\varepsilon > 0$, define

$$P_n \equiv \Pr\left(\text{for some } \ell,\ \left| \sum_{j=\ell+1}^{\ell+r_n} (Y_j - 1) \right| > r_n \varepsilon \right),$$

and note that

$$P_n \leq \sum_{\ell=1}^{n-r_n} \Pr\left(\sum_{j=\ell+1}^{\ell+r_n} (Y_j - 1) > r_n\varepsilon\right) + \sum_{\ell=1}^{n-r_n} \Pr\left(\sum_{j=\ell+1}^{\ell+r_n} (Y_j - 1) < -r_n\varepsilon\right). \tag{14}$$

We need to show that the right hand side goes to zero. Using the Chernoff bound (based on Markov's inequality), which says that $\Pr(X > 0) \leq E\left[e^{tX}\right]$ for any random variable X and $t > 0$ such that the right side is finite, we obtain

$$\Pr\left(\sum_{j=\ell+1}^{\ell+r_n} (Y_j - 1) > r_n\varepsilon\right) \leq E\left[e^{t\left(\sum_{j=\ell+1}^{\ell+r_n}(Y_j-1)-r_n\varepsilon\right)}\right] = E\left[e^{t\left(\sum_{j=\ell+1}^{\ell+r_n} Y_j - r_n - r_n\varepsilon\right)}\right].$$

Letting $Z \equiv \sum_{j=\ell+1}^{\ell+r_n} Y_j$, we have

$$\Pr\left(\sum_{j=\ell+1}^{\ell+r_n} (Y_j - 1) > r_n\varepsilon\right) \leq E\left[e^{t(Z-r_n-r_n\varepsilon)}\right] = E\left[e^{tZ}e^{-tr_n(1+\varepsilon)}\right] = e^{-tr_n(1+\varepsilon)}E\left[e^{tZ}\right].$$

Because the sum of r_n exponential random variables with mean 1 has the Gamma distribution with shape parameter r_n and scale parameter 1, Z is a Gamma$(r_n, 1)$ random variable. Because the moment generating function for a Gamma$(r_n, 1)$ random variable is $M(t) =$

$\frac{1}{(1-t)^{r_n}}$ for $0 < t < 1$, it follows that for $0 < t < 1$, $E[e^{tZ}] = \frac{1}{(1-t)^{r_n}}$. Thus, for all $0 < t < 1$,

$$\Pr\left(\sum_{j=\ell+1}^{\ell+r_n}(Y_j - 1) > r_n\varepsilon\right) \le e^{-tr_n(1+\varepsilon)}\frac{1}{(1-t)^{r_n}} = \left(\frac{e^{-t(1+\varepsilon)}}{1-t}\right)^{r_n}.$$

Choosing the minimizing value $t^* = \frac{\varepsilon}{1+\varepsilon}$ gives the bound

$$\Pr\left(\sum_{j=\ell+1}^{\ell+r_n}(Y_j - 1) > r_n\varepsilon\right) \le \left(\frac{e^{-t^*(1+\varepsilon)}}{1-t^*}\right)^{r_n} = \left((1+\varepsilon)e^{-\varepsilon}\right)^{r_n}.$$

A similar bound holds for each term of the second sum on the right side of (14), and there are $n - r_n$ terms in each sum. Therefore,

$$P_n \le 2(n-r_n)\left((1+\varepsilon)e^{-\varepsilon}\right)^{r_n} = 2(n-r_n)\left(\frac{e^{\varepsilon}}{1+\varepsilon}\right)^{-r_n} = 2(n-r_n)a(\varepsilon)^{-r_n},$$

where $a(\varepsilon) > 1$ for $\varepsilon > 0$. We need to show that $2(n-r_n)a(\varepsilon)^{-r_n} \to 0$. To do this, note that our assumption that $\frac{\ln n}{r_n} \to 0$ implies that $\frac{1}{nk'(n)} \to 0$, so

$$\begin{aligned}
\lim_{n\to\infty}\frac{2(n-r_n)}{a(\varepsilon)^{r_n}} &\le \lim_{n\to\infty}\frac{2n}{a(\varepsilon)^{r_n}} \\
&= \frac{2}{\ln a(\varepsilon)}\lim_{n\to\infty}\frac{1}{a(\varepsilon)^{r_n}k'(n)} \\
&\le \frac{2}{\ln a(\varepsilon)}\lim_{n\to\infty}\frac{1}{nk'(n)} \\
&= 0,
\end{aligned}$$

where the first equality uses L'Hôpital's rule and the second inequality uses $\frac{n}{a^{r_n}} \to 0$ (which follows from L'Hôpital's rule), which implies that $a(\varepsilon)^{r_n} > n$ for n sufficiently large. The final equality uses $\frac{1}{nk'(n)} \to 0$. Thus, we have shown that $P_n \to 0$, and so (13) is proved, which proves (11). This completes Part 1 of the proof.

Part 2: To conclude the proof, we show that

$$\Pr\left[\sup_{\ell\in\{1,\dots,n-r_n\}}\left|\frac{U_{r_n}(V_{(\ell)})}{d_{r_n}(V_{(\ell)})} - f(V_{(\ell)})\right| > \varepsilon\right] \to 0. \tag{15}$$

It follows from (11) that $\Pr\left[\sup_{\ell\in\{1,\dots,n-r_n\}}\left|U_{r_n}(V_{(\ell)})\right| > \varepsilon\right] \to 0$ and hence, because f is everywhere positive, that

$$\Pr\left[\sup_{\ell\in\{1,\dots,n-r_n\}} \left|d_{r_n}(V_{(\ell)})\right| > \varepsilon\right] \to 0. \tag{16}$$

Note that

$$\begin{aligned}
\left|\frac{U_{r_n}(V_{(\ell)})}{d_{r_n}(V_{(\ell)})} - f(V_{(\ell)})\right| &= \left|\frac{F(V_{(\ell)}) - F(V_{(\ell)} - d_{r_n}(V_{(\ell)}))}{d_{r_n}(V_{(\ell)})} - f(V_{(\ell)})\right| \\
&= \left|\frac{1}{d_{r_n}(V_{(\ell)})}\int_{V_{(\ell)}-d_{r_n}(V_{(\ell)})}^{V_{(\ell)}} \big(f(t) - f(V_{(\ell)})\big)\, dt\right| \\
&\le \max_{t\in[V_{(\ell)}-d_{r_n}(V_{(\ell)}),V_{(\ell)}]} \left|f(t) - f(V_{(\ell)})\right|,
\end{aligned}$$

where the first equality uses the definition of U_{r_n}, the second equality uses the differentiability of F. Using (16) and the uniform continuity of f (f is continuous on a compact support and so uniformly continuous), (15) follows.

Conclusion: Thus, we have shown that

$$\Pr\left(\sup_{\ell\in\{1,\dots,n-r_n\}} \left|\frac{n}{r_n}U_{r_n}(V_{(\ell)}) - 1\right| > \varepsilon\right) \to 0 \text{ and } \Pr\left(\sup_{\ell\in\{1,\dots,n-r_n\}} \left|\frac{U_{r_n}(V_{(\ell)})}{d_{r_n}(V_{(\ell)})} - f(V_{(\ell)})\right| > \varepsilon\right) \to 0.$$

Putting these together, we have

$$\Pr\left(\sup_{\ell\in\{1,\dots,n-r_n\}} \left|\frac{1}{\frac{n}{r_n}U_{r_n}(V_{(\ell)})}\frac{U_{r_n}(V_{(\ell)})}{d_{r_n}(V_{(\ell)})} - f(V_{(\ell)})\right| > \varepsilon\right) \to 0,$$

which, rearranging, gives us

$$\Pr\left(\sup_{\ell\in\{1,\dots,n-r_n\}} \left|\frac{r_n}{n}\frac{1}{d_{r_n}(V_{(\ell)})} - f(V_{(\ell)})\right| > \varepsilon\right) \to 0,$$

which completes the proof because $\frac{r_n}{n}\frac{1}{d_{r_n}(V_{(\ell)})} = \frac{1}{n\sigma^v_\ell} = \hat{f}(V_{(j)})$, which is our implied density estimator. ■

References

Akbarpour, M. and S. Li (2018): "Credible Mechanisms," 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 (2016): "Group Strategy-Proofness in Private Good Economies," *American Economic Review*, 106, 1073–99.

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

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 (2019): "Two-Sided Allocation Problems, Decomposability, and the Impossibility of Efficient Trade," *Journal of Economic Theory*, 179, 416–454.

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.

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 (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," *American Economic Review*, 107, 3257–3287.

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 L. M. Marx (2019): "Merger Review for Markets with Buyer Power," *Journal of Political Economy*, 127.

Loertscher, S. and C. Mezzetti (2018): "A Dominant Strategy, Double Clock Auction with Estimation-Based Tâtonnement," Working Paper, University of Melbourne.

Loftsgaarden, D. O. and C. P. Quesenberry (1965): "Nonparametric Estimate of a Multivariate Density Function," *Annals of Mathematical Statistics*, 36, 1049–1051.

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 (forth.): "Deferred-Acceptance Clock Auctions and Radio Spectrum Reallocation," *Journal of Political Economy.*

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

Moore, D. and E. Henrichon (1969): "Uniform Consistency of Some Estimates of a Density Function," *Annals of Mathematical Statistics*, 40, 1499–1502.

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.

Nagaraja, H. N., K. Bharath, and F. Zhang (2015): "Spacings Around an Order Statistic," *Annals of the Institute of Statistical Mathematics*, 67, 515–540.

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.