Available online at www.sciencedirect.com

ScienceDirect

Journal of Economic Theory 187 (2020) 105030

JOURNAL OF Economic Theory

www.elsevier.com/locate/jet

Asymptotically optimal prior-free clock auctions ☆

Simon Loertscher [a], Leslie M. Marx [b,*]

[a] *Department of Economics, Level 4, FBE Building, 111 Barry Street, University of Melbourne, Victoria 3010, Australia*
[b] *Fuqua School of Business, Duke University, 100 Fuqua Drive, Durham, NC 27708, USA*

Received 16 January 2019; final version received 3 March 2020; accepted 15 March 2020
Available online 20 March 2020

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.

JEL classification: C72; D44; L13

Keywords: Asymptotic optimality; Estimating virtual types; Spacings

☆ We thank Alessandro Pavan (the editor), three anonymous referees, Salvador Barberà, Arthur Campbell, 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, from the Faculty of Business and Economics at the University of Melbourne via an Eminent Research Scholar Grant, and from the Australian Research Council under a Discovery Project Grant DP200103574. Bing Liu provided outstanding research assistance.

* Corresponding author.
E-mail addresses: simonl@unimelb.edu.au (S. Loertscher), marx@duke.edu (L.M. Marx).

https://doi.org/10.1016/j.jet.2020.105030

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 the agent's virtual type, but also on the agent's ranking relative to the other agents on the same side of the market, which is determined using the bids of the other players on that 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 the agent's report to determine other agents' virtual types and their relative rankings may indirectly introduce a means to manipulate the mechanism.[4]

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

[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 (2020).

[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 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.

allocation.[5] Assuming that agents play their obviously dominant strategies and, on each side of the market, draw their types independently 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 mechanism 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 (2020) 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 homogeneous within groups but heterogeneous across groups, we can allow for price discrimination across groups, revenue thresholds, group-specific caps, and group-specific favoritism.

While for the purpose of compactness and generality, we present our prior-free clock auction as a double clock auction in which buyers and sellers have to be given incentives to reveal decision-relevant information, it equally applies to auction settings in which only one side of the market has to be incentivized. As such, our paper provides an explanation for the widely observed phenomena of dynamically adjusting, secret reserves (without claiming exclusivity of such an explanation). Within the domain of two-sided settings, it provides a dynamic price discovery process that may prove useful for platforms and intermediaries that match buyers and sellers. Importantly, it does so without relying on discriminatory treatment of agents, which is not only problematic when the designer is a government, but also for some private entities. More broadly and fundamentally, we provide a dynamic allocation process that allows the allocation decisions to be based on information about type distributions that is gleaned from agents who participate in the process without distorting any agent's incentives. In our setting, this type of information is valuable because the designer is interested in generating revenue. However, in dynamic settings in which agents arrive sequentially, even social surplus maximization requires such distributional information.[6]

[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 for sufficiently large numbers of buyers and sellers and there is a unique local maximum under price posting.

[6] As an example, consider the problem of assigning cadaver kidneys to patients on a wait list. Both patients and organs arrive stochastically over time, and the optimal assignment will not only depend on the pool of patients present, but also on the composition of the pools of patients and organs that will arrive in the future.

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 (2020), and Li (2017), this literature has identified advantages of dynamic implementation over direct mechanisms in a variety of setups.[7] In particular, our paper builds on the properties of clock auctions identified by Milgrom and Segal (2020) and on design features first introduced by McAfee (1992).[8]

Motivated by Wilson (1987), we develop prior-free clock auctions that are asymptotically optimal in the sense that, as the number of traders goes to infinity, the ratio of the profit of the designer from using the prior-free clock auction over the profit it would obtain if it knew the traders' distributions and applied the optimal mechanism derived by Myerson (1981), goes to one.[9] For expositional purposes, we focus on two-sided exchanges as do Myerson and Satterthwaite (1983), Gresik and Satterthwaite (1989), and Williams (1999). However, analogous results hold for the case of one-sided private information, as we show in Appendix A. 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.[10]

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).[11] 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

[7] Although both the English auction and the second-price auction make bidding truthfully a dominant strategy, in laboratory settings subjects are consistently more likely to play their dominant strategy in English auctions than in second-price auctions (Kagel, 1995), suggesting that the open format of the English auction facilitates discovery of the dominant strategies. As is well known, with single-unit demands, the English auction is a dynamic implementation of the second-price auction that uses the exit points of agents who have become inactive as the price for the remaining active bidders; for more elaborate analyses of setups with one-sided private information in which the designer can survey losing bidders and has an incentive to use information obtained in that way to set the reserve, see the working papers by Brooks (2013a,b).

[8] 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.

[9] 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.

[10] 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.

[11] 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 the 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).

the distribution of types for other agents, referred to as random-sampling mechanisms.[12] Whereas the analysis of this type of mechanism in Baliga and Vohra (2003) focuses on profit maximization for the designer, the literature on Algorithmic Game Theory focuses on whether the mechanisms have good worst-case performance relative to benchmarks based on prior-free mechanisms that approximate Bayesian optimality but are not incentive compatible.[13] 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.[14] None of these random-sampling mechanisms can be implemented as a clock auction or is envy-free.

The same is true for the simple variant of a random-sampling mechanism in which a subset of agents is randomly selected as the set of agents whose reports are used to estimate distributions and those agents are prevented from trading and make and receive no payments. See, for example, Baliga and Vohra (2003) and Segal (2003) for a discussion of this type of mechanism. While such mechanisms are natural from a statistical perspective, their violation of envy-freeness and non-capriciousness is a severe problem not only for designers that are government agents, and are often prohibited by law from discriminating among otherwise identical agents,[15] but also for private companies that are keen to at least appear nondiscriminatory in their pricing.[16] Moreover, as noted by Baliga and Vohra (2003, Remark 3) and Segal (2003, p. 510), the informational inefficiency associated with prohibiting a subset of agents from trading slows convergence and results in worse per-capita performance.

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 itself? The short answer is that it 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 our setup with two-sided private information, Bayesian optimality, and clock auctions. Among other things, it describes 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

[12] 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).

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

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

[15] For example, the U.S. Administrative Procedure Act (5 USC §706(2)(A)) requires that courts "hold unlawful and set aside agency action, findings, and conclusions found to be," among other things, "arbitrary" or "capricious." As an example, the spectrum license auction procedures of the U.S. Federal Communications Commission are subject to court challenges arguing that they are arbitrary and/or capricious.

[16] See, for example, Chen et al. (2015), who write that as of April 2015, Uber's "surge prices were no longer uniform for users, even if they were in the same surge area at the same time. We reported this finding to Uber, and their engineers confirmed that this behavior was caused by a bug in their system."

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 the paper. We provide details for the setup with one-sided private information in Appendix A and longer proofs in Appendix B.

2. Setup

In the body of this paper, we study a two-sided setting in which there is a demand side composed of buyers with unit demands and a supply side composed of sellers with unit capacities. We provide results for the corresponding one-sided settings in Appendix A. We let $\mathcal{N}$ denote the set of n buyers and $\mathcal{M}$ denote the set of m sellers. Buyers are privately informed of their values, and sellers are privately informed of their costs. We let $\mathbf{v}$ denote the vector of the n buyers' values and $\mathbf{c}$ denote the vector of the m sellers' marginal costs. 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 trades occur via a monopoly market maker, who is a risk-neutral designer without private information.

We assume that buyers draw their types independently from the same distribution F with continuous density f that is positive on the support $[\underline{v}, \overline{v}]$.[17] Similarly, we assume that sellers draw their types independently from the same distribution G with continuous density g that is positive on the support $[\underline{c}, \overline{c}]$. In the online appendix, we extend our results to allow for heterogeneity across groups of buyers and heterogeneity across groups of sellers.[18]

When considering prior-free mechanisms, we assume that the mechanism designer knows neither the agents' types nor the distributions from which they are drawn, but does know 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. In addition, we assume that the designer knows upper and lower bounds for the agents' 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. Further, we assume that the buyers and sellers do not know the distributions from which their types are drawn. Keeping fixed the mechanism, equilibrium behavior would not be affected if the agents did know the distributions because the mechanism endows them with dominant strategies; however, in that case, depending on assumptions about the informational structure, alternative mechanisms, such as those developed by Crémer and McLean (1985, 1988), could be optimal.[19]

In the prior-free setting with private types that we consider, where neither the designer nor any of the agents knows the type distributions, dominant-strategy incentive compatibility is the natural notion of incentive compatibility, and ex post individual rationality is the appropriate

[17] Our asymptotic results rely on the continuity of the inverses of the virtual type functions (defined below), which is guaranteed if the densities are continuous.

[18] Specifically, in the extension we assume that all buyers within a group draw their values independently from the same distribution and that all sellers within a group draw their costs independently 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.

[19] 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.

notion of individual rationality.[20] Because the agents and the designer lack a (common) prior, we would simply not know how to impose incentive-compatibility and individual-rationality constraints in any other, meaningful way because we would not know how to predict agents' behavior.[21] In our setting, these notions of incentive compatibility and individual rationality thus stem from what may be called a common lack of statistical information.[22]

2.1. Bayesian optimality with two-sided private information

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) as the *no-full trade* condition, and throughout the paper we assume that the no-full trade condition holds. Under this condition and the assumption that the virtual value function Φ and the virtual cost function Γ, which are defined as

$$\Phi(v) \equiv v - \frac{1 - F(v)}{f(v)} \quad \text{and} \quad \Gamma(c) \equiv c + \frac{G(c)}{g(c)},$$

are increasing, the Bayesian optimal mechanism is characterized by the allocation rule that, given $(\mathbf{v}, \mathbf{c})$, induces q trades, where q is the largest index such that $\Phi(v_{(q)}) \geq \Gamma(c_{[q]})$.[23] 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.

We expand the scope of the analysis by assuming that, in the Bayesian benchmark in which the designer knows the distributions F and G, the designer wants to maximize a Ramsey objective, that is, a weighted sum of expected profit and social surplus,[24] 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 rule as for the case of profit maximization, except that one has to replace the virtual value and virtual cost functions by, respectively, the relevant *weighted virtual type function* $\Phi_\alpha(v)$ and $\Gamma_\alpha(c)$ defined as

---

20 With private values, dominant strategy incentive compatibility is equivalent to ex post incentive compatibility, so we could, of course, equivalently impose ex post incentive compatibility.

21 Although our designer does not know the distributions, the designer does know that this is an independent privates values environment. Because of the well-known 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)), the focus on dominant strategy incentive compatibility and ex post individual rationality is without loss of generality once the designer has learned the relevant aspects of the distributions, which in our asymptotically optimal clock auction, the designer will learn.

22 For foundations of dominant-strategy implementation based on the designer's preferences, see Chung and Ely (2007) and Yamashita (2015) for single-object allocation problems and Chen and Li (2018) for more general settings.

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

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

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

By construction, for $\alpha = 0$, the weighted virtual types correspond to true types, so that $\alpha = 0$ is equivalent to ex post efficiency. Moreover, because $\Phi(v) = \Phi_1(v)$ and $\Gamma(c) = \Gamma_1(c)$, the case with $\alpha = 1$ corresponds to profit maximization.[25] Under the assumption that Φ and Γ are increasing, it follows that for all $\alpha \in [0, 1]$, Φ_α and Γ_α are also increasing. The payments in the dominant-strategy implementation are accordingly defined by replacing the virtual type functions Φ and Γ (and their inverses) by Φ_α and Γ_α (and their inverses) in the formulas above.

Beyond generality, allowing the designer to have a Ramsey objective also highlights the need for estimation because, as soon as $\alpha > 0$, the optimal mechanism depends on distributional details. Similarly, as noted in the introduction, in dynamic settings social surplus maximization is no longer a prior-free exercise, which provides additional motivation for the development of dynamic mechanisms that allow the designer or planner to glean distributional information from participating agents without distorting their incentives.

2.2. Asymptotic optimality

Now that we have identified the Bayesian optimal outcome, we can define what it means for a mechanism to be asymptotically optimal. We follow Gresik and Satterthwaite (1989) in their approach to deriving asymptotic results by considering replicas of the economy, whereby a given initial set of agents is replicated, but each replicated agent has an independently drawn type. This structure ensures that the numbers of buyers and sellers go to infinity in fixed proportion.

Fixing the initial number of buyers n and sellers m, each replica adds an additional n buyers and m sellers, with each buyer independently drawing its value from F and each seller independently drawing its cost from G. Given $\eta \in \{1, 2, ...\}$, an η-fold replica of this initial economy has ηn buyers and ηm sellers, each with an independently drawn type. For all replicas, the designer's beliefs about the type space amount to beliefs over F and G, which is the same as for a single replica. We say that an incentive compatible, individually rational mechanism defined for an η-fold replica is *asymptotically optimal* if the ratio of the expected value of the designer's objective in the mechanism relative to the expected value of the designer's objective in the Bayesian optimal mechanism converges to 1 in probability as the number of replicas goes to infinity.

2.3. Clock auctions

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

[25] 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-multiplier 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 $\alpha = \lambda/(1+\lambda)$.

Loosely speaking, in a clock auction, active buyers and sellers choose whether to exit as the buyer clock price increases and the seller clock price decreases. 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. We assume that each buyer observes at least the buyer clock price and that 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 its 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 its cost.

To formally define a clock auction for our setup, we adapt the definition of a clock auction in Milgrom and Segal (2020) to accommodate a two-sided setting.[26] 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 a clock auction is defined by the functions ϕ, γ, and τ, we denote a clock auction by $\mathcal{C}_{\phi,\gamma,\tau}$.

In words, the clock auction proceeds as follows: 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 designer checks the condition for the termination of the auction, which is whether the value of the buyer function is greater than or equal to the value of the seller function, i.e., $\phi(\boldsymbol{\omega}) \geq \gamma(\boldsymbol{\omega})$. If it is satisfied, then the auction ends; otherwise, the designer uses the target function $\tau(\boldsymbol{\omega})$ to determine target prices for both the buyer and seller clocks and advances the clocks until either both target prices are reached with no exits, in which case the auction ends, or an exit occurs, in which case the procedure repeats, starting with the equalization of the number of active buyers and sellers.

In more detail, 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. Specifically, $\mathcal{M}^A \subseteq \mathcal{M}$ is the set of active sellers, m^A is the cardinality of $\mathcal{M}^A$, $\mathbf{x}^S \in \mathbb{R}^{m-m^A}$ is the vector of exit prices for sellers who exited, and $p^S \in \mathbb{R}$ is the seller clock price.

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.

[26] Milgrom and Segal (2020) 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.

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 $\boldsymbol{\omega}_{t+1} = (1, \boldsymbol{\omega}_t^B, \boldsymbol{\omega}_t^S)$ and the auction ends. 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 $\overline{p}$,[27] 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 $\underline{p}$,[28] 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 $z_{t+1} = 1$.

By the usual logic, playing truthful strategies is dominant-strategy incentive compatible in a clock auction.

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 $\phi(\boldsymbol{\omega}) = p^B$ and $\gamma(\boldsymbol{\omega}) = p^S$, and the target function gives the midpoint between the two clock prices, $\tau(\boldsymbol{\omega}) = \frac{p^B + p^S}{2}$.

Key properties of clock auctions

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

For environments like ours 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 the buyer's value, then the buyer's payoff is zero, and if a

27 Under our no-full trade condition and truthful bidding, this does not occur.

28 As in the case above, under our no-full trade condition and truthful bidding, this does not occur.

buyer remains active after the buyer clock price reaches the buyer's value, then the buyer's payoff is bounded above by zero; however, under truthful bidding, the buyer's 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 its 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 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 whether the Bayesian optimal mechanism can be implemented by a clock auction.

Proposition 1. *With two-sided private information, the Bayesian optimal mechanism cannot be implemented by a clock auction.*

Proof. Milgrom and Segal (2020, Proposition 7) show that a necessary condition for clock implementation is that agents are substitutes. However, when private information pertains to both sides of the market, buyers and sellers are complements. To see this, note that the problem is assignment representable, as defined by Delacrétaz et al. (2019), and so the complementarity of buyers and sellers follows from Shapley (1962). (To apply the results of Delacrétaz et al. (2019) and Shapley (1962), which focus on ex post efficiency, to our setting, simply replace the social surplus that is created when buyer i is matched to seller j, $v_i - c_j$, by the corresponding virtual surplus, $\Phi_\alpha(v_i) - \Gamma_\alpha(c_j)$.) ■

Proposition 1 summarizes the implication of the results of Milgrom and Segal (2020) for our setting. 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. In contrast, as we describe in Appendix A, in settings with one-sided private information, the Bayesian optimal

[29] On weak group strategy-proofness in a one-sided clock auction, see Li (2017) and Milgrom and Segal (2020). On the connection between individual and group strategy-proofness, see Barberà et al. (2016). 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.

mechanism can be defined based only on the information held by the designer and information gleaned from nontrading agents, and thus has a clock auction implementation.

3. Asymptotically optimal clock auctions

We now show the existence of a prior-free clock auction that is asymptotically optimal. To have a prior-free clock auction $\mathcal{C}_{\phi,\gamma,\tau}$, we require a buyer function ϕ, seller function γ, and target function τ, all of 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 further constraints on how these functions are estimated. In addition, we show that the structure of the buyer and seller functions can 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.

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

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

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

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

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 $c_{[j]}$. Endowed with consistent estimates of the virtual type functions, the allocation based on the intersection of these estimated virtual types will then quite naturally converge to the allocation rule that is optimal when the type distributions (and hence the virtual type functions) are known.

Pursuing this idea, we establish the asymptotic optimality of the prior-free clock auction $\mathcal{C}_{\phi,\gamma,\tau}$ defined below. We ignore the target function τ 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 consider the buyer function ϕ and seller function γ given by

$$\phi(\boldsymbol{\omega}) = p^B - \alpha j \sigma_j^v \text{ and } \gamma(\boldsymbol{\omega}) = p^S + \alpha j \sigma_j^c, \tag{2}$$

where σ_j^v and σ_j^c are spacing estimators.

To obtain consistent estimators of the expected spacing between values $v_{(j-1)}$ and $v_{(j)}$ and between the costs $c_{[j]}$ and $c_{[j-1]}$ based on values $v_{(j)}, \ldots, v_{(n)}$ and costs $c_{[j]}, \ldots, c_{[m]}$, we use the average of $r(n)$ and $r(m)$ spacings for nearby worse types, where $r : \mathbb{R}_+ \to \mathbb{R}_+$ is assumed to be differentiable and to satisfy:

$$\lim_{j\to\infty} r(j) = \infty \quad \text{and} \quad \lim_{j\to\infty} \frac{r(j)}{j} = 0 \quad \text{and} \quad \lim_{j\to\infty} \frac{\ln j}{r(j)} = 0. \tag{3}$$

For example, the conditions in (3) are satisfied if $r(j) = j^x$ for $x \in (0, 1)$. These conditions ensure that $r'(j)$ goes to zero as j goes to infinity; however, to make sure that it does not go to zero too fast, we require, via the final condition in (3), that $\frac{1}{jr'(j)}$ goes to zero, i.e., $jr'(j)$ goes to infinity, as j goes to infinity.[31] Loosely, we need $r(j)$ to shrink relative to j fast enough, but not too fast.

Given exit prices for buyers of $\hat{v}_{(n)}, \ldots, \hat{v}_{(j)}$ and exit prices for sellers of $\hat{c}_{[m]}, \ldots, \hat{c}_{[j]}$ (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 $\hat{v}_{(j)}$, divided by $r(n)$, and similarly for σ_j^c. More precisely, to account for the possibility that $r(n)$ or $r(m)$ exceeds the number of available observations, we define the spacing estimators as follows:

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

As we now show, $j\sigma_j^v$ provides a uniformly consistent estimator of $\frac{1-F(v_{(j)})}{f(v_{(j)})}$ and $j\sigma_j^c$ provides a uniformly consistent estimator of $\frac{G(c_{[j]})}{g(c_{[j]})}$. 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 $F(v_{(j)})$, which is $(n-j)/n$, then we have $j\sigma_j^v = \frac{1-(n-j)/n}{\hat{f}(v_{(j)})} = \frac{j/n}{\hat{f}(v_{(j)})}$, 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 of the density f when evaluated at $v_{(j)}$.

By the Glivenko-Cantelli Theorem (van der Vaart, 1998, Theorem 19.1), 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, together with the assumption that the densities are positive, completes the proof that $j\sigma_j^v$ is a uniformly consistent estimator of $\frac{1-F(v_{(j)})}{f(v_{(j)})}$ and, analogously, that $j\sigma_j^c$ is a uniformly consistent estimator of $\frac{G(c_{[j]})}{g(c_{[j]})}$.

[31] To see that $\frac{1}{jr'(j)}$ goes to zero as j goes to infinity, notice that, by the first condition in (3), $r(j)$ goes to infinity, as does, of course, $\ln j$. Applying L'Hôpital's Rule, we obtain $\lim_{j\to\infty} \frac{\ln j}{r(j)} = \frac{1}{jr'(j)}$, which is zero by the last condition in (3).

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. *Recalling that we assume that F and G are continuously differentiable with positive densities f and g and compact supports and that r is differentiable and satisfies* (3)*, then as* $n \to \infty$*,*

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

and as $m \to \infty$*,*

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

Sketch of proof and intuition. Here we provide a sketch of the proof and offer some intuition. Appendix B provides the full proof. We focus on the distribution of values because the result for costs is analogous.

Given $V_1, ..., V_n$ iid random draws from F, define $d_{r(n)}(z)$ to be the distance from z to the $r(n)$-th closest of the random draws that are less than or equal to z. Thus, $d_{r(n)}(V_{(\ell)}) = r(n)\sigma_\ell^v$. As mentioned above, the density estimator implied by our virtual value estimator is

$$\hat{f}(V_{(\ell)}) = \frac{1}{n\sigma_\ell^v} = \frac{r(n)}{n}\frac{1}{d_{r(n)}(V_{(\ell)})}.$$

The proof shows that this estimator converges uniformly (for all $\ell \in \{1, ..., n-r(n)\}$) in probability to $f(V_{(\ell)})$.

The proof proceeds by showing that as n goes to infinity, first, $\frac{F(V_{(\ell)})-F(V_{(\ell)}-d_{r(n)}(V_{(\ell)}))}{r(n)/n}$ converges uniformly in probability to 1 and, second, $\frac{F(V_{(\ell)})-F(V_{(\ell)}-d_{r(n)}(V_{(\ell)}))}{d_{r(n)}(V_{(\ell)})}$ converges uniformly in probability to $f(V_{(\ell)})$. Then, putting these two together, we have the result that $\frac{r(n)}{n}\frac{1}{d_{r(n)}(V_{(\ell)})}$ converges uniformly in probability to $f(V_{(\ell)})$, which completes the proof.

The second step is relatively straightforward using the definition of a density, our differentiability assumptions,[32] and the result that $d_{r(n)}(V_{(\ell)})$ converges uniformly to zero, which relies on the assumptions that $\lim_{n\to\infty}\frac{r(n)}{n} = 0$ and that f is positive everywhere on the support. Thus, for the remainder of this sketch, we focus on the first step of the proof.

For the first step of the proof, define $U_{r(n)}(z) \equiv F(z) - F(z - d_{r(n)}(z))$ and note that because the distance $d_{r(n)}(V_{(\ell)})$ is made up of distances between $r(n)$ pairs of adjacent order statistics, we have:

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

[32] Our assumption that F is continuously differentiable with a compact support implies that f is continuous on a compact support and so uniformly continuous, which is required for this result.

We need to show that $\frac{n}{r(n)} U_{r(n)}(V_{(\ell)})$ converges uniformly in probability to 1. The demonstration has two main components.

In the first component, we use a result relating the values of a cdf at adjacent order statistics to exponential random variables (see, e.g., Nagaraja et al., 2015) to conclude that $U_{r(n)}(V_{(\ell)})$ has the same distribution as $\sum_{j=\ell+1}^{\ell+r(n)} \frac{Y_j}{S_{n+1}}$, where $\{Y_j\}_{j=1}^{n+1}$ are independent exponential random variables with mean 1 and $S_{n+1} \equiv \sum_{j=1}^{n+1} Y_j$.

In the second component, we show that $\frac{n}{r(n)} \sum_{j=\ell+1}^{\ell+r(n)} \frac{Y_j}{S_{n+1}}$ converges uniformly in probability to 1. To develop some intuition for why one might expect this to hold, note that $U_{r(n)}(V_{(\ell)})$ has the same asymptotic properties as the sum of $r(n)$ out of $n+1$ mean-1 random variables divided by the sum of all $n+1$ random variables. Therefore, loosely speaking, the sum of the $r(n)$ mean-1 random variables in the numerator "cancels" with the $r(n)$ in the denominator, and the n in the numerator "cancels" with the sum of $n+1$ mean-1 random variables in the denominator. Thus, the ratio stays close to 1 for n sufficiently large.

More precisely, the result follows because (i) S_{n+1} is the sum of $n+1$ iid random variables with mean 1, and so by the law of large numbers $\frac{S_{n+1}}{n}$ converges almost surely to 1 and (ii) the sums $\sum_{j=\ell+1}^{\ell+r(n)} \frac{Y_j}{r(n)}$ are uniformly near 1 in probability. This last bit takes some work to show and requires the use of (i) the Chernoff bound (based on Markov's inequality), (ii) the result that the sum of $r(n)$ exponential random variables with mean 1 has a Gamma$(r(n), 1)$ distribution, (iii) our assumption that $r(n) \to \infty$, and (iv) our assumption that $\frac{\ln n}{r(n)} \to 0$, which implies that $\frac{1}{nr'(n)} \to 0$, which using L'Hôpital's rule allows us to pin down one of the required limits. This completes the proof that $\frac{n}{r(n)} U_{r(n)}(V_{(\ell)})$ converges uniformly in probability to 1, and together with the other elements, then completes the proof. □

Using Theorem 1, and assuming that (3) is satisfied, it follows that $v_{(j)} - j\sigma_j^v$ and $c_{[j]} + j\sigma_j^c$ 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 $\mathcal{C}_{\phi,\gamma,\tau}$ is asymptotically optimal if ϕ and γ satisfy* (2)–(4).

Proof. See Appendix B.

The foundation for Proposition 2 is the result that as the number of agents goes to infinity, the estimated virtual types defined by (2)–(4) 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. The proof of Proposition 2 formalizes this intuition and shows that the differences in per-trader payments between the prior-free clock auction and the optimal mechanism converge in probability to zero as the number of replicas goes to infinity. The latter is essentially a consequence of the payoff-equivalence theorem and the fact that the

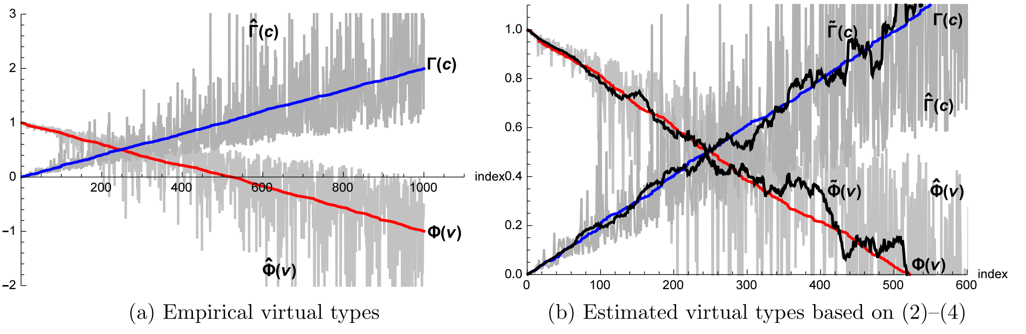

Fig. 1. Panel (a): Theoretical virtual types Φ and Γ and local empirical virtual types $\hat{\Phi}$ and $\hat{\Gamma}$. Panel (b): Theoretical virtual types Φ and Γ, local empirical virtual types $\hat{\Phi}$ and $\hat{\Gamma}$, and estimated virtual types $\tilde{\Phi}(v_{(j)}) \equiv v_{(j)} - j\sigma_j^v$ and $\tilde{\Gamma}(c_{[j]}) \equiv c_{[j]} + j\sigma_j^c$ with $r(j) = j^{1/2}$. Both panels are based on 1000 values and 1000 costs drawn from the uniform distribution on [0, 1].

allocation in the prior-free clock auction converges in probability to the allocation in the optimal mechanism.

Our estimated virtual type functions differ from the "empirical" virtual type functions in that we use a uniformly convergent (in probability) estimator of 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 Fig. 1 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—that is, with r satisfying (3)—is asymptotically optimal.

Fig. 1 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 (2)–(4) provide a much better approximation of the theoretical virtual types in that example and deliver an approximately optimal number of trades.

Using similar figures, we can illustrate that, as long as there is a unique optimum (e.g., if the virtual type functions are monotone), then for large numbers of agents, discontinuities do not create difficulties for our estimator, except perhaps resulting some number of "lost" trades while the estimated virtual types adjust to the discontinuity. As an illustration, in the example shown in Fig. 2, the theoretical optimum occurs at the point of discontinuity with approximately 600 of the 1000 pairs of agents trading, but in the prior-free clock auction based on estimated virtual types result, there are only approximately 575 trades—a loss of $25/600 = 4\%$ of trades relative to the optimum.

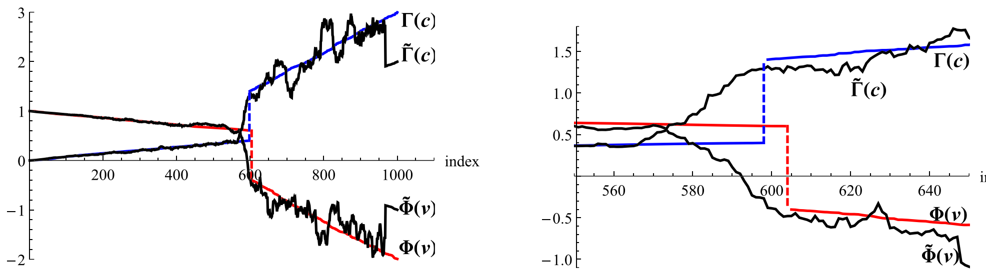

(a) Estimated virtual types

(b) Estimated virtual types: close-up

Fig. 2. Panel (a): Theoretical virtual types Φ and Γ and estimated virtual types $\tilde{\Phi}(v_{(j)}) \equiv v_{(j)} - j\sigma_j^v$ and $\tilde{\Gamma}(c_{[j]}) \equiv c_{[j]} + j\sigma_j^c$ with $r(j) = j^{1/2}$. Panel (b): Close-up view of Panel (a). Both panels are based on 1000 values drawn from F and 1000 costs drawn from G, where the distributions have support $[0, 1]$ and $F(v) = v/2$ for $v \in [0, 4/5]$ and $F(v) = 3v - 2$ for $v \in (4/5, 1]$ and $G(c) = 3c$ for $c \in [0, 1/5]$ and $G(c) = x/2 + 1/2$ for $c \in (1/5, 1]$.

3.2. Sequential consistency

In this subsection, we consider whether there is additional structure that can be brought to bear on the functional form of the virtual type estimators. As shown in Proposition 2, the prior-free clock auction defined in (2)–(4) is asymptotically optimal. However, the asymptotic optimality result continues to hold for any specification of ϕ and γ of the form:

$$\phi(\boldsymbol{\omega}) = p^B - \chi_{\alpha,j}\sigma_j^v \text{ and } \gamma(\boldsymbol{\omega}) = p^S + \chi_{\alpha,j}\sigma_j^c, \tag{5}$$

where $\chi_{\alpha,j}$ is nonnegative (to ensure that the mechanism is deficit free) and satisfies the condition that $\lim_{j\to\infty} \chi_{\alpha,j}/j = \alpha$, which ensures that $\chi_{\alpha,j}$ has the asymptotic properties of αj, which is the relevant condition for Theorem 1 to hold. Thus, the virtual type estimators are, so far, only pinned down by the asymptotic properties of the coefficient $\chi_{\alpha,j}$. However, as we now discuss, other considerations besides asymptotics allow us to put additional structure on that coefficient.

Consider the thought experiment in which the designer delegates the operation of a dynamic mechanism to a decision maker, such as an auctioneer, and suppose that the auctioneer's objective is aligned with the designer's in that it puts weight α on revenue and weight $1 - \alpha$ on efficiency, but, consistent with our desire for a deficit-free mechanism, with a payoff of $-\infty$ if the clock auction ends with a deficit. We ask what would need to be true of the coefficient $\chi_{\alpha,j}$ for it to be credible that the auctioneer would follow the protocol defined by our prior-free clock auction. In the prior-free clock auction, key decision points occur when there has been an exit that leaves the auction in a state with equal numbers of buyers and sellers. Suppose that at such points, the auctioneer can choose either to end the auction or to continue the auction. (We focus on decision point following an exit, avoiding the additional complexities associated with how an auctioneer might interact with target prices specified by the mechanism.)

The expected profitability of stopping or continuing the auction depends on the auctioneer's expectations regarding the types of the "next" buyer and seller to exit. Thus, in order to analyze the auctioneer's incentives, we need to specify the auctioneer's expectations. To put structure on the problem, we assume that the auctioneer's expectations when $j - 1$ active buyers and sellers remain are given by the estimated spacings σ_j^v and σ_j^c. Specifically, we assume that the auctioneer's expectation of $v_{(j-1)}$ is that it is equal to $v_{(j)} + \sigma_j^v$ and that the auctioneer's expectation of $c_{[j-1]}$ is that it is equal to $c_{[j]} - \sigma_j^c$. Given these expectations, we ask whether the auctioneer's incentives at decision points following an exit are consistent with the rules of the prior-free clock auction.

The notion we have in mind is similar to the credibility notion developed independently by Akbarpour and Li (2020), but 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]

We say that a prior-free clock auction satisfying (4) and (5) is *sequentially consistent* if for all $j \in \{1, ..., \min\{m, n\}\}$ and $\boldsymbol{\omega} \in \Omega_j$, the buyer and seller functions call for the end of the auction, i.e., $\phi(\boldsymbol{\omega}) \geq \gamma(\boldsymbol{\omega})$, if and only if ending the auction would not result in a deficit, i.e., $p^B \geq p^S$, and ending the auction is consistent with the auctioneer's objective and expectations, i.e.,

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

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

$$v_{(j)} - (\alpha(j-2) - (1-\alpha))\sigma_j^v \geq c_{[j]} + (\alpha(j-2) - (1-\alpha))\sigma_j^c. \tag{6}$$

Because we are considering a state in which the clock prices are defined by the exits of the j-th highest valuing buyer and j-th lowest cost seller, it follows that $p^B = v_{(j)}$ and $p^S = c_{[j]}$. Thus, both $p^B \geq p^S$ and (6) hold if and only if

$$p^B - \chi_{\alpha,j}\sigma_j^v \geq p^S + \chi_{\alpha,j}\sigma_j^c, \tag{7}$$

where

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

giving us the following result:

Proposition 3. *The prior-free clock auction $\mathcal{C}_{\phi,\gamma,\tau}$ with ϕ and γ satisfying (3)–(5) is sequentially consistent if and only if $\chi_{\alpha,j}$ satisfies (8).*

It follows from Proposition 3, that for an asymptotically optimal, sequentially consistent, prior-free clock auction $\mathcal{C}_{\phi,\gamma,\tau}$ with a designer's objective given by α, one would use the buyer function

$$\phi(\boldsymbol{\omega}) = p^B - \max\{0, \alpha(j-2) - (1-\alpha)\}\sigma_j^v$$

and seller function

$$\gamma(\boldsymbol{\omega}) = p^S - \max\{0, \alpha(j-2) - (1-\alpha)\}\sigma_j^c,$$

[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.

where r satisfies (3) (see Section 3.4 for discussion) and τ is a feasible target function (see Section 3.3 for discussion).

3.3. Target functions

The asymptotic optimality result for the prior-free clock auction $\mathcal{C}_{\phi,\gamma\,\tau}$ described above did not require that we specify the target function τ because τ affects the number of trades by at most one and so is not relevant for asymptotic results. However, the role of the target function can be significant with small numbers of traders. Thus, in this section, we provide a methodology for defining the target function.

For a state $\boldsymbol{\omega}$ with $j-1$ active buyers and sellers, one can define $\tau(\boldsymbol{\omega}) \in [\phi(\boldsymbol{\omega}), \gamma(\boldsymbol{\omega})]$ to maximize the probability of $j-1$ trades when there should be $j-1$ trades, i.e., to maximize 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 function" 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 function is given by:

$$\tau(\boldsymbol{\omega}) = \min\left\{\gamma(\boldsymbol{\omega}),\ \max\left\{\phi(\boldsymbol{\omega}),\ \frac{\phi(\boldsymbol{\omega}) + \gamma(\boldsymbol{\omega})}{2} + \left(1 - \frac{\alpha}{2}\right)\left(\sigma_j^v - \sigma_j^c\right)\right\}\right\}. \tag{9}$$

According to (9), 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. For example, as mentioned above, McAfee (1992) focuses on efficiency, i.e., the case with $\alpha = 0$. Indeed, the mechanism of McAfee (1992) is the clock auction $\mathcal{C}_{\phi,\gamma,\tau}$ with ϕ and γ defined by (5) and (8) with $\alpha = 0$ and with $\tau(\boldsymbol{\omega}) = \frac{\phi(\boldsymbol{\omega})+\gamma(\boldsymbol{\omega})}{2} = \frac{p^B+p^S}{2}$, which corresponds to the target function in (9), but with no attempt at estimation, i.e., with $\sigma_j^v = \sigma_j^c = 0$.

3.4. 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 satisfies (3).

Further, for a sample of size n, 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. The numerical results displayed there also demonstrate that convergence is fast, with 90 percent or more of the profit that would be obtained using the optimal mechanism being achieved by the prior-free clock auction when there are only single-digit numbers of buyer-seller pairs.

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

3.5. Discussion

In this subsection, we discuss regularity assumptions and the calculation of confidence intervals for the estimated gain from continuing an auction.

Regularity assumptions

The Bayesian optimal mechanism described above rests on the assumption of increasing virtual type functions, which is what, following Myerson (1981), has become known 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 as higher values (lower costs) imply higher virtual values (lower virtual costs). 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 $\Gamma(c') < \overline{v}$,[34] then the Bayesian optimal mechanism differs from the one described above. For finite numbers of buyers and sellers, random rationing occurs with positive probability because the point of intersection between (ironed) virtual values and (ironed) virtual costs, which determines the quantity traded, occurs with positive probability in the "flat" (or ironed) parts of these functions.

For our purposes, we can relax the usual regularity condition considerably. To illustrate, 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 is that the prices satisfy

$$\Phi(p^B) = \Gamma(p^S) \text{ and } n(1 - F(p^B)) = mG(p^S). \tag{10}$$

If the solution to (10) is unique and if the Bayesian optimal mechanism with sufficiently large numbers of buyers and sellers 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 for sufficiently large numbers of agents even when they are not Bayesian optimal with a small number of traders (see Fig. 3 for an example).[35] Correspondingly, our results for asymptotically optimal prior-free mechanisms require only the assumption that there is a unique pair of prices satisfying (10) and that the Bayesian optimal mechanism for sufficiently large numbers of agents is a posted-price mechanism. In what follows, we maintain this assumption.

The tight connection between the Bayesian optimality of posted-price mechanism with sufficiently large numbers of agents and the asymptotic optimality of our prior-free clock auction is not surprising; after all, 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 (10) 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 (10) holds, with Φ and Γ replaced by prior-free estimates.

---

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

[35] When private information only pertains to buyers, respectively sellers, the optimal posted prices are obtained from (10) by replacing Γ and Φ by the identity function.

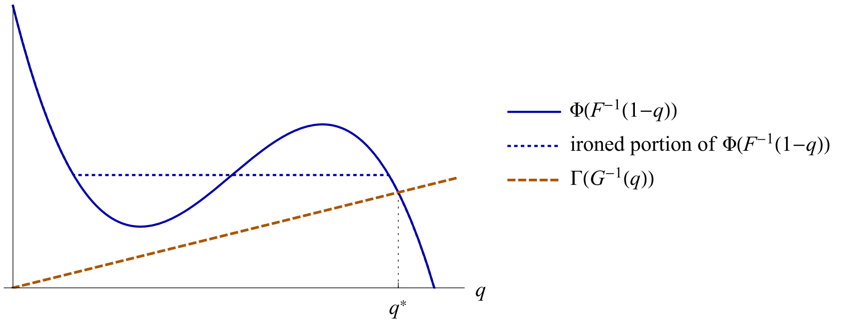

Fig. 3. Illustration of a setup with unique posted prices that are Bayesian optimal for sufficiently large numbers of agents, but where ironing may be required in the small.

As an illustration, Fig. 3 depicts a case in which, with large, 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, with small numbers of traders, 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 with sufficiently large numbers of agents would no longer be a posted-price mechanism because some share of agents would have to be rationed.

Confidence intervals

We now illustrate briefly how the clock auction can also be augmented to generate confidence intervals that provide additional guidance to the designer (or the auctioneer). For the purpose of this illustration, we assume $\alpha = 1$. 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).[36] This is illustrated in Fig. 4 for an example with 20 buyers and 20 sellers.[37] Panel (a) shows the estimated 95% confidence

[36] 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.

[37] As an alternative, for certain choices of r, 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 in Section 3.4 is required (up to proportionality) for minimizing mean square error.

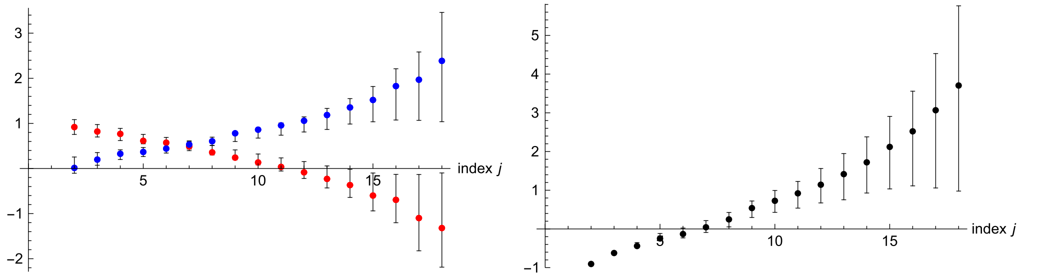

(a) Estimated virtual types with CIs (b) Estimated revenue changes with CIs

Fig. 4. 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 $n = m = 20$ and $r(j) = j^{4/5}$, with values and costs drawn from the uniform distribution on [0,1].

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 Fig. 4(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 following this exit. Turning to Fig. 4(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.

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 heterogeneous 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.

Second-degree price discrimination

When the agents in different buyer and seller groups draw their types independently from different distributions, the Bayesian optimal mechanism still induces trade among the sellers with the lowest virtual costs and the buyers with the highest virtual types. But now there is an element of second-degree price discrimination on each side of the market because agents in different groups are characterized by different virtual type functions. 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. In the online appendix, we show how our analysis and clock auction generalize to such settings.

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 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 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. For such an extension when the designer's objective is efficiency, see Loertscher and Marx (2020).

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 the mechanism performs, a designer employing an asymptotically optimal mechanism facing many traders may find it reassuring to know that it would not have chosen any other mechanism at the outset had it then known the distributions that it has inferred now. In that sense, asymptotic optimality, like privacy preservation, protects the designer from regret.

Appendix A. One-sided private information

In this appendix, we provide the details for the setups with one-side private information pertaining to buyers and pertaining to sellers.

Bayesian optimality with one-sided private information pertaining to buyers

For the setup with *one-sided private information pertaining to buyers*, we let $c_{[1]}, ..., c_{[m]}$ be the commonly known marginal cost curve of the designer. The n buyers have unit demands and draw their valuations independently from distribution F with continuous density f that is positive on the support $[\underline{v}, \overline{v}]$. A buyer's payoff is equal to the buyer's value minus the price that the buyer pays if the buyer trades and zero otherwise. It is well known that under the no-full trade assumption that $\underline{v} \leq c_{[1]} < \overline{v}$ and the assumption that the virtual valuation function Φ 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 $\Phi(v_{(q)}) \geq c_{[q]}$.[38] The q buyers with the highest values trade and, in the dominant-strategy implementation, pay the price $p^B = \max\{v_{(q+1)}, \Phi^{-1}(c_{[q]})\}$.[39] 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 distribution G with continuous density g that is positive on the support $[\underline{c}, \overline{c}]$. A seller's payoff is equal to the payment that the seller receives minus the seller's cost if the seller trades and zero otherwise. If the designer's marginal values $v_{(1)}, ..., v_{(n)}$ are commonly known, the no-full trade condition holds, i.e., $\overline{c} \geq v_{(1)} > \underline{c}$, and the virtual cost function Γ is increasing, then this is a standard procurement auction in which the optimal quantity traded is the largest index q such that $v_{(q)} \geq$

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

[39] 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.

$\Gamma(c_{[q]})$.[40] The q sellers with the lowest costs trade and, in the dominant-strategy implementation, are paid $p^S = \min\{c_{[q+1]}, \Gamma^{-1}(v_{(q)})\}$.[41]

Clock auction for a one-sided setup

The 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.

Results for one-sided private information

In contrast to the negative result in Proposition 1 for two-sided private information, with one-sided private information, the Bayesian optimal mechanism can be implemented by a clock auction:

Proposition A.1. *Under the assumption of 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.*

Proof. In our setting, if the virtual type function is increasing, then Milgrom and Segal (2020, Proposition 7) implies that clock implementation is possible when agents are substitutes. The result then follows once one notices that, with private information pertaining to only one side of the market, the privately informed agents are indeed substitutes for each other. ■

As an illustration of Proposition A.1, 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, $\phi(\boldsymbol{\omega}) = \Phi_\alpha(p_B)$, $\gamma(\boldsymbol{\omega}) = c_{[n^A+1]}$, and $\tau(\boldsymbol{\omega}) = c_{[n^A]}$.

The adaptation of Proposition 2 on the asymptotic optimality of our prior-free clock auction 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 function need not be estimated. For example, with private information only on the buyer side, if the state $\boldsymbol{\omega}$ has $j-1$ active buyers, then the seller function $\gamma(\boldsymbol{\omega})$ and the target function $\tau(\boldsymbol{\omega})$ would both be replaced by $c_{[j-1]}$.

Proposition A.2. *In the setup with one-sided private information pertaining to buyers, the one-sided prior-free clock auction with ϕ and σ_j^v satisfying* (2)–(4) *is asymptotically optimal; simi-*

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

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

larly, with private information pertaining to sellers, the one-sided prior-free clock auction with γ and σ_j^c satisfying (2)–(4) *is asymptotically optimal.*

In addition, Proposition 3 related to the sequential consistency of the prior-free clock auction continues to hold in the one-sided setup.

Appendix B. 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

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

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 $n \to \infty$,

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

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 $\ell\in\{1,\ldots,n-r(n)\}$,

$$\begin{aligned}\sum_{j=1}^{r(n)}\left(F(V_{(\ell+j-1)})-F(V_{(\ell+j)})\right)&=F(V_{(\ell)})-F(V_{(\ell+r(n))})\\&=F(V_{(\ell)})-F(V_{(\ell)}-d_{r(n)}(V_{(\ell)}))=U_{r(n)}(V_{(\ell)}).\end{aligned}$$

Of course, then

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

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,

$$F(V_{(n)}),\quad F(V_{(n-1)})-F(V_{(n)}),\quad \ldots,\quad F(V_{(1)})-F(V_{(2)}),\quad 1-F(V_{(1)})$$

have the same joint distribution as

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

where $Y_1,\ldots,Y_{n+1}$ are independent exponential random variables with mean 1 and $S_{n+1}=Y_1+\ldots+Y_{n+1}$. 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 $\ell\in\{1,\ldots,n-r(n)\}$) in probability to 1, i.e., that for all $\varepsilon>0$,

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

which we can rewrite as

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

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 $\ell\in\{1,\ldots,n-r(n)\}$) 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

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

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

$$\begin{aligned}\Pr\left(\sum_{j=\ell+1}^{\ell+r(n)} (Y_j - 1) > r(n)\varepsilon\right) &\leq \mathbb{E}\left[e^{t\left(\sum_{j=\ell+1}^{\ell+r(n)}(Y_j-1)-r(n)\varepsilon\right)}\right] \\ &= \mathbb{E}\left[e^{t\left(\sum_{j=\ell+1}^{\ell+r(n)} Y_j - r(n) - r(n)\varepsilon\right)}\right].\end{aligned}$$

Letting $Z \equiv \sum_{j=\ell+1}^{\ell+r(n)} Y_j$, we have

$$\begin{aligned}\Pr\left(\sum_{j=\ell+1}^{\ell+r(n)} (Y_j - 1) > r(n)\varepsilon\right) &\leq \mathbb{E}\left[e^{t(Z-r(n)-r(n)\varepsilon)}\right] \\ &= \mathbb{E}\left[e^{tZ}e^{-tr(n)(1+\varepsilon)}\right] = e^{-tr(n)(1+\varepsilon)}\mathbb{E}\left[e^{tZ}\right].\end{aligned}$$

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

$$\Pr\left(\sum_{j=\ell+1}^{\ell+r(n)} (Y_j - 1) > r(n)\varepsilon\right) \leq e^{-tr(n)(1+\varepsilon)} \frac{1}{(1-t)^{r(n)}} = \left(\frac{e^{-t(1+\varepsilon)}}{1-t}\right)^{r(n)}.$$

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

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

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,

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

where $a(\varepsilon) \equiv \frac{e^{\varepsilon}}{1+\varepsilon}$, so $a(\varepsilon) > 1$ for $\varepsilon > 0$. We need to show that $2(n - r(n))(a(\varepsilon))^{-r(n)} \to 0$. To do this, recall that our assumption that $\frac{\ln n}{r(n)} \to 0$ implies by L'Hôpital's rule that $\frac{1}{nr'(n)} \to 0$, so

$$\begin{aligned}\lim_{n\to\infty} \frac{2(n - r(n))}{(a(\varepsilon))^{r(n)}} &\leq \lim_{n\to\infty} \frac{2n}{(a(\varepsilon))^{r(n)}} \\ &= \frac{2}{\ln a(\varepsilon)} \lim_{n\to\infty} \frac{1}{(a(\varepsilon))^{r(n)} r'(n)}\end{aligned}$$

$$\leq \frac{2}{\ln a(\varepsilon)} \lim_{n\to\infty} \frac{1}{nr'(n)}$$
$$= 0,$$

where the first equality uses L'Hôpital's rule and the second inequality uses $\frac{n}{(a(\varepsilon))^{r(n)}} \to 0$ (which follows from L'Hôpital's rule and $r(n) \to \infty$), which implies that $(a(\varepsilon))^{r(n)} > n$ for n sufficiently large. The final equality uses $\frac{1}{nr'(n)} \to 0$. 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

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

It follows from (11) and the assumption that $\lim_{n\to\infty} \frac{r(n)}{n} = 0$ that

$$\Pr\left[\sup_{\ell\in\{1,\ldots,n-r(n)\}} \left|U_{r(n)}(V_{(\ell)})\right| > \varepsilon\right] \to 0$$

and hence, because f is assumed everywhere positive, that

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

Note that

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

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

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

Putting these together, we have

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

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 $\frac{r(n)}{n}\frac{1}{d_{r(n)}(V_{(\ell)})}=\frac{1}{n\sigma_\ell^v}=\hat{f}(V_{(j)})$, which is our implied density estimator. ■

Proof of Proposition 2. From Theorem 1 we know that the estimated virtual types converge in probability to the true virtual types in any prior-free clock auctions $\mathcal{C}_{\phi,\gamma,\tau}$ in which ϕ and γ satisfy (2)–(4). To prove Proposition 2, we are left to show, first, that the allocation rule of the prior-free clock auction that is based on these estimates converges in probability to the allocation rule of the optimal mechanism. Then, second, we show that the ratio of the profit from using $\mathcal{C}_{\phi,\gamma,\tau}$ to the profit from the optimal mechanism converges in probability to 1 by invoking the payoff-equivalence theorem and continuity to show that the probability that the per-buyer profits and the per-seller expenditures differ goes to zero as the number of replicas goes to infinity.

Recall that we consider replicas of a market with n buyers and m sellers, where the η-th replica has $n\eta$ buyers and $m\eta$ sellers. Define v^* and c^* by $n(1-F(v^*))=mG(c^*)$ and $\Phi(v^*)=\Gamma(c^*)$. Note that in the limit as η goes to infinity, the Bayesian optimal mechanism involves trade by all buyers with value greater than v^*, i.e., share $1-F(v^*)$ of the $n\eta$ buyers, and trade by all sellers with cost less than c^*, i.e., share $G(c^*)$ of the $m\eta$ sellers. Note also that for all $\varepsilon>0$,

$$\Phi(v^*-\varepsilon)<\Phi(v^*)=\Gamma\left(G^{-1}\left[\frac{n}{m}(1-F(v^*))\right]\right)<\Gamma\left(G^{-1}\left[\frac{n}{m}(1-F(v^*-\varepsilon))\right]\right) \quad (17)$$

and

$$\Phi(v^*+\varepsilon)>\Phi(v^*)=\Gamma\left(G^{-1}\left[\frac{n}{m}(1-F(v^*))\right]\right)>\Gamma\left(G^{-1}\left[\frac{n}{m}(1-F(v^*+\varepsilon))\right]\right). \quad (18)$$

Let $\mathcal{C}_{\phi,\gamma,\tau}$ be a prior-free clock auction where ϕ and γ satisfy (2)–(4). Given n buyers and m sellers, let $\tilde{\Phi}_n(v_{(j)})\equiv v_{(j)}-j\sigma_j^v$ and $\tilde{\Gamma}_m(c_{[j]})\equiv c_{[j]}+j\sigma_j^c$. Let $\tilde{v}_{n,m}$ be the random variable such that when there are n buyers and m sellers, all buyers with value greater than $\tilde{v}_{n,m}$ trade in $\mathcal{C}_{\phi,\gamma,\tau}$. By definition, $\tilde{v}_{n,m}$ satisfies $\tilde{\Phi}_n(\tilde{v}_{n,m})\geq\tilde{\Gamma}_m(\tilde{c}_{n,m})$, where $\tilde{c}_{n,m}$ is defined to be such that the number of sellers with cost less than $\tilde{c}_{n,m}$ is equal to the number of buyers with value greater than $\tilde{v}_{n,m}$.

Let $\varepsilon>0$ be given and suppose that

$$\lim_{\eta\to\infty}\Pr\left(\tilde{v}_{n\eta,m\eta}<v^*-\varepsilon\right)>0. \quad (19)$$

This implies that

$$\lim_{\eta\to\infty}\Pr\left(\tilde{\Phi}_{n\eta}(v^*-\varepsilon)\geq\tilde{\Gamma}_{m\eta}\left(G^{-1}\left[\frac{n}{m}(1-F(v^*-\varepsilon))\right]\right)\right)>0. \quad (20)$$

But, using Theorem 1, this implies that

$$\Phi(v^*-\varepsilon)\geq\Gamma\left(G^{-1}\left[\frac{n}{m}(1-F(v^*-\varepsilon))\right]\right). \quad (21)$$

To see this, note that if (21) does not hold, then, letting $\delta\equiv\Gamma\left(G^{-1}\left[\frac{n}{m}(1-F(v^*-\varepsilon))\right]\right)-\Phi(v^*-\varepsilon)>0$, Theorem 1 implies that

$$\lim_{\eta\to\infty}\Pr\left(\tilde{\Gamma}_{m\eta}\left(G^{-1}\left[\frac{n}{m}(1-F(v^*-\varepsilon))\right]\right)-\tilde{\Phi}_{n\eta}(v^*-\varepsilon)>\delta/2\right)=1,$$

which contradicts (20). But (21) contradicts (17), allowing us to conclude that our supposition in (19) does not hold. Analogously, let $\varepsilon > 0$ be given and suppose that

$$\lim_{\eta\to\infty} \Pr\left(\tilde{v}_{n\eta,m\eta} > v^* + \varepsilon\right) > 0. \tag{22}$$

This implies that

$$\lim_{\eta\to\infty} \Pr\left(\tilde{\Phi}_{n\eta}(v^*+\varepsilon) \le \tilde{\Gamma}_{m\eta}\left(G^{-1}\left[\frac{n}{m}(1-F(v^*+\varepsilon))\right]\right)\right) > 0.$$

But, as above, using Theorem 1, this implies that

$$\Phi(v^*+\varepsilon) \le \Gamma\left(G^{-1}\left[\frac{n}{m}(1-F(v^*+\varepsilon))\right]\right),$$

which contradicts (18). Thus, we conclude that for all $\varepsilon > 0$,

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

A similar contradiction arises if we do the analysis based on $\tilde{c}_{n\eta,m\eta}$ differing from c^*, which allows us to conclude that

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

In other words, as $\eta \to \infty$, the probability that the thresholds for buyer and seller types that trade under the prior-free clock auction $\mathcal{C}_{\phi,\gamma,\tau}$ and under the optimal mechanism differ by more than ε, for any $\varepsilon > 0$, is zero.

Given cutoffs $\tilde{v}_{\eta n,\eta m}$ and v^* under the clock auction and the optimal mechanism, the difference in the per-buyer expected payments is, by the payoff-equivalence theorem,

$$\left|\int_{\tilde{v}_{\eta n,\eta m}}^{v^*} \Phi(v) f(v) dv\right|,$$

and similarly the difference in per-seller expected payment given cutoffs $\tilde{c}_{\eta n,\eta m}$ and c^* is

$$\left|\int_{\tilde{c}_{\eta n,\eta m}}^{c^*} \Gamma(c) g(c) dc\right|.$$

We are left to show that these differences converge in probability to zero. Fix $\Delta > 0$ and focus on the difference in per-buyer expected payments (the argument for per-seller payments is analogous). We need to show that

$$\lim_{\eta\to\infty} \Pr\left(\left|\int_{\tilde{v}_{\eta n,\eta m}}^{v^*} \Phi(v) f(v) dv\right| > \Delta\right) = 0.$$

To show this, let $\tilde{v} = \tilde{v}_{\eta n,\eta m}$ for ease of notation and note that for any $\varepsilon > 0$,

$$\Pr\left(\left|\int_{\tilde{v}}^{v^*}\Phi(v)f(v)dv\right|>\Delta\right)$$

$$=\Pr\left(\left|\int_{\tilde{v}}^{v^*}\Phi(v)f(v)dv\right|>\Delta \mid |\tilde{v}-v^*|\geq\varepsilon\right)\Pr\left(|\tilde{v}-v^*|\geq\varepsilon\right)$$

$$+\Pr\left(\left|\int_{\tilde{v}}^{v^*}\Phi(v)f(v)dv\right|>\Delta \mid |\tilde{v}-v^*|<\varepsilon\right)\Pr\left(|\tilde{v}-v^*|<\varepsilon\right)$$

$$\leq\Pr\left(|\tilde{v}-v^*|\geq\varepsilon\right)+\Pr\left(\left|\int_{\tilde{v}}^{v^*}\Phi(v)f(v)dv\right|>\Delta \mid |\tilde{v}-v^*|<\varepsilon\right).$$

The first term in the last line of the expression above is zero in the limit as η goes to infinity by the preceding arguments, and because $\Phi(v)$ and $f(v)$ are bounded, the second term is also zero in the limit as η goes to infinity as we can make ε arbitrarily small. Thus,

$$\lim_{\eta\to\infty}\Pr\left(\left|\int_{\tilde{v}}^{v^*}\Phi(v)f(v)dv\right|>\Delta\right)=0$$

as required. ■

Appendix C. Supplementary material

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.jet.2020.105030.

References

Akbarpour, M., Li, S., 2020. Credible auctions: A trilemma. Econometrica 88, 425–467.
Arrow, K., 1959. Toward a theory of price adjustment. In: Abramovitz, M. (Ed.), The Allocation of Economic Resources: Essays in Honor of Bernard Francis Haley. Stanford University Press, Stanford, pp. 41–51.
Athey, S., Haile, P.A., 2007. Nonparametric approaches to auctions. In: Heckman, J.J., Leamer, E.E. (Eds.), Handbook of Econometrics, vol. 6A. Elsevier, pp. 3847–3966.
Ausubel, L., 2004. An efficient ascending-bid auction for multiple objects. Am. Econ. Rev. 94, 1452–1475.
Ausubel, L., 2006. An efficient dynamic auction for heterogeneous commodities. Am. Econ. Rev. 96, 602–629.
Baliga, S., Vohra, R., 2003. Market research and market design. Adv. Theor. Econ. 3, 5.
Barberà, S., Berga, D., Moreno, B., 2016. Group strategy-proofness in private good economies. Am. Econ. Rev. 106, 1073–1099.
Brooks, B.A., 2013a. Revenue Sharing in Second-Price Auctions. University of Chicago.
Brooks, B.A., 2013b. Surveying and Selling: Belief and Surplus Extraction in Auctions. University of Chicago.
Bulow, J., Roberts, J., 1989. The simple economics of optimal auctions. J. Polit. Econ. 97, 1060–1090.
Chen, L., Mislove, A., Wilson, C., 2015. Peeking beneath the hood of Uber. In: MC '15 Proceedings of the 2015 Internet Measurement Conference, pp. 495–508.
Chen, Y.-C., Li, J., 2018. Revisiting the foundations of dominant-strategy mechanisms. J. Econ. Theory 178, 294–317.
Chung, K.-S., Ely, J., 2007. Foundations of dominant strategy mechanisms. Rev. Econ. Stud. 74, 447–476.
Crémer, J., McLean, R., 1985. Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent. Econometrica 53, 345–361.

Crémer, J., McLean, R., 1988. Extraction of the surplus in Bayesian and dominant strategy auctions. Econometrica 56, 1247–1257.

Cripps, M.W., Swinkels, J.M., 2006. Efficiency of large double auctions. Econometrica 74, 47–92.

Dasgupta, P., Maskin, E., 2000. Efficient auctions. Q. J. Econ. 115, 341–388.

Delacrétaz, D., Loertscher, S., Marx, L.M., Wilkening, T., 2019. Two-sided allocation problems, decomposability, and the impossibility of efficient trade. J. Econ. Theory 179, 416–454.

Deshmukh, K., Goldberg, A.V., Hartline, J.D., Karlin, A.R., 2002. Truthful and competitive double auctions. In: Proceedings of the 10th European Symposium on Algorithms. Springer-Verlag.

Devanur, N.R., Hartline, J.D., Yan, Q., 2015. Envy freedom and prior-free mechanism design. J. Econ. Theory 156, 103–143.

Dhangwatnotai, P., Roughgarden, T., Yan, Q., 2015. Revenue maximization with a single sample. Games Econ. Behav. 91, 318–333.

Dütting, P., Roughgarden, T., Talgam-Cohen, I., 2017. Modularity and greed in double auctions. Games Econ. Behav. 405, 59–83.

Gershkov, A., Goeree, J.K., Kushnir, A., Moldovanu, B., Shi, X., 2013. On the equivalence of Bayesian and dominant strategy implementation. Econometrica 81, 197–220.

Goldberg, A.V., Hartline, J.D., Karlin, A.R., Saks, M., Wright, A., 2006. Competitive auctions. Games Econ. Behav. 55, 242–269.

Goldberg, A.V., Hartline, J.D., Wright, A., 2001. Competitive auctions and digital goods. In: Proceedings of the Twelfth Annual ACM-SIAM Symposium on Discrete Algorithms, ACM/SIAM, pp. 735–744.

Gresik, T., Satterthwaite, M., 1989. The rate at which a simple market converges to efficiency as the number of traders increases: an asymptotic result for optimal trading mechanisms. J. Econ. Theory 48, 304–332.

Guerre, E., Perrigne, I., Vuong, Q., 2000. Optimal nonparametric estimation of first-price auctions. Econometrica 68, 525–574.

Hagerty, K., Rogerson, W., 1987. Robust trading mechanisms. J. Econ. Theory 42, 94–107.

Hurwicz, L., 1973. The design of mechanisms for resource allocation. Am. Econ. Rev. 63, 1–30.

Kagel, J.H., 1995. Auctions: a survey of experimental research. In: Kagel, J.H., Roth, A.E. (Eds.), The Handbook of Experimental Economics. Princeton University Press, Princeton, NJ, pp. 501–585.

Levin, J., Skrzypacz, A., 2016. Properties of the combinatorial clock auction. Am. Econ. Rev. 106, 2528–2551.

Li, S., 2017. Obviously strategy-proof mechanisms. Am. Econ. Rev. 107, 3257–3287.

Loertscher, S., Marx, L., Wilkening, T., 2015. A long way coming: Designing centralized markets with privately informed buyers and sellers. J. Econ. Lit. 53, 857–897.

Loertscher, S., Marx, L.M., 2020. A dominant-strategy asset market mechanism. Games Econ. Behav. 120, 1–15.

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

Loftsgaarden, D.O., Quesenberry, C.P., 1965. Nonparametric estimate of a multivariate density function. Ann. Math. Stat. 36, 1049–1051.

Lopomo, G., 1998. The English auction is optimal among simple sequential auctions. J. Econ. Theory 82, 144–166.

Lopomo, G., 2000. Optimality and robustness of the English auction. Games Econ. Behav. 36, 219–240.

Lucking-Reiley, D., 2000. Vickrey auctions in practice: from nineteenth-century philately to twenty-first-century E-commerce. J. Econ. Perspect. 14, 183–192.

Manelli, A.M., Vincent, D.R., 2010. Bayesian and dominant-strategy implementation in the independent private-values model. Econometrica 78, 1905–1938.

McAdams, D., Schwarz, M., 2007. Credible sales mechanisms and intermediaries. Am. Econ. Rev. 97, 260–276.

McAfee, R.P., 1992. A dominant strategy double auction. J. Econ. Theory 56, 434–450.

McMillan, J., 1994. Selling spectrum rights. J. Econ. Perspect. 8, 145–162.

Milgrom, P., 2004. Putting Auction Theory to Work. Cambridge University Press.

Milgrom, P., Segal, I., 2020. Clock auctions and radio spectrum reallocation. J. Polit. Econ. 128, 1–31.

Milgrom, P., Weber, R., 1982. A theory of auctions and competitive bidding. Econometrica 50, 1089–1122.

Moore, D., Henrichon, E., 1969. Uniform consistency of some estimates of a density function. Ann. Math. Stat. 40, 1499–1502.

Moore, D.S., Yackel, J.W., 1977. Large sample properties of nearest neighbor density function estimators. In: Gupta, S.S., Moore, D.S. (Eds.), Statistical Decision Theory and Related Topics II. Academic Press, pp. 269–279.

Myerson, R., Satterthwaite, M., 1983. Efficient mechanisms for bilateral trading. J. Econ. Theory 29, 265–281.

Myerson, R.B., 1981. Optimal auction design. Math. Oper. Res. 6, 58–73.

Nagaraja, H.N., Bharath, K., Zhang, F., 2015. Spacings around an order statistic. Ann. Inst. Stat. Math. 67, 515–540.

Norman, P., 2004. Efficient mechanisms for public goods with use exclusions. Rev. Econ. Stud. 71, 1163–1188.
Rustichini, A., Satterthwaite, M.A., Williams, S.R., 1994. Convergence to efficiency in a simple market with incomplete information. Econometrica 62, 1041–1063.
Satterthwaite, M., Williams, S., Zachariadis, K., 2015. Price Discovery Using a Double-Auction. SSRN Working Paper.
Segal, I., 2003. Optimal pricing mechanisms with unknown demand. Am. Econ. Rev. 93, 509–529.
Shapley, L., 1962. Complements and substitutes in the optimal assignment problem. Nav. Res. Logist. Q. 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. J. Finance 16, 8–37.
Williams, S.R., 1999. A characterization of efficient, Bayesian incentive compatible mechanisms. Econ. Theory 14, 155–180.
Wilson, R., 1987. Game-theoretic analyses of trading processes. In: Bewley, T.F. (Ed.), Advances in Economic Theory, Fifth World Congress. Cambridge University Press, pp. 33–70.
Yamashita, T., 2015. Implementation in weakly undominated strategies: optimality of second-price auction and posted-price mechanism. Rev. Econ. Stud. 82, 1223–1246.