Prior-free Bayesian Optimal Double-Clock Auctions*

Simon Loertscher† Leslie M. Marx‡

November 12, 2015

Abstract

We develop a prior-free double-clock auction that is asymptotically Bayesian optimal. It preserves the privacy of trading buyers and sellers, endows agents with obviously dominant strategies, is envy-free and weakly group strategy-proof. We characterize the Bayesian optimal privacy preserving mechanism, show that our mechanism converges to it as estimation errors vanish, and establish the impossibility of ex post efficient, or indeed optimal, privacy-preserving trade. Our mechanism accommodates revenue thresholds and other constraints, requires limited commitment by the designer, extends to accommodate heterogeneous groups of agents, permits real-time diagnostics, and performs well in the small.

Keywords: price discovery, privacy preservation, two-sided mechanisms, revenue maximization, estimating virtual types, spacing
JEL Classification: C72, D44, L13

---

*We thank Salvador Barberà, Matt Jackson, Claudio Mezzetti, Andras Niedermayer, Stefan Reichelstein, Mark Satterthwaite, Ilya Segal, Thomas Tröger, Steve Williams, Robert Wolpert, and seminar participants at the University of Melbourne, University of Rochester, University of Technology in Sydney, the 2015 Decentralization Conference at Northwestern, the 2015 MaCCI Summer Institute, and the 2015 EARIE Conference in Munich for helpful comments.

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

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

1 Introduction

Bayesian mechanism design has proved fruitful as an analytical tool and a conceptual framework, yet has been deemed "fragile" for practical purposes on the grounds that it depends on the fine details of the environment about which designers, and possibly the agents, are likely to be uncertain. This has led to the postulates that for the purposes of practicality, mechanisms ought to be free of the details of the environment, such as the distributions from which agents draw their types, and robust with respect to agents' beliefs about the environment. At the same time, and to the extent that econometricians and savvy designers can infer distributions and other design-relevant details, the theory of Bayesian mechanism design is empirically predictive exactly because the mechanisms it prescribes depend on the fine details of the environment. As a case in point, in an optimal auction, the reserve price varies with the distributions from which bidders draw their values. This raises the question as to whether it is possible to design mechanisms that are both practical, in the sense of being robust and prior free, and empirically predictive insofar as the mechanism varies with the environment.

In this paper, we answer this question in the affirmative by developing a double-clock auction that is prior free, endows agents with dominant strategies, and converges to the Bayesian optimal mechanism as the number of agents becomes large, while still being operational and envy free for any size of the market.[1] We do so in a two-sided environment with privately informed buyers and sellers. Although our double-clock auction is defined without reference to distributions and does not depend on the agents' or designer's beliefs about distributions, its equilibrium outcomes vary with distributions because the mechanism estimates the relevant details and uses these estimates to determine who trades. The double-clock auction preserves the privacy of trading agents and endows agents with obviously dominant strategies.[2] It is constructed in such a way that it requires only limited commitment by the designer, and it permits real-time diagnostics as to whether revenue is expected to increase or decrease if the auction continues. Our mechanism is weakly group strategy-proof because of the VCG nature of the pricing rule, and the resulting equilib-

[1]A mechanism is envy free if no agent wishes to swap allocations and payments with another agent. Not only is envy-freeness normatively appealing and sometimes even imposed by law, but it is also necessary for approximate implementation in large economies as shown by Jackson and Kremer (2007).

[2]The concept of "obviously dominant strategies" is introduced by Li (2015). The English auction endows bidders with obviously dominant strategies because at the earliest point in the game at which a strategy differs from the dominant strategy, the maximum payoff from the alternative strategy is not more than the minimum payoff a bidder gets from playing the dominant strategy. In contrast, the second-price auction does not endow agents with obviously dominant strategies because the minimum payoff from bidding truthfully is zero whereas when overbidding a buyer may win at a price below her value. The experimental evidence summarized by Kagel (1995) supports the view that experimental subjects are more likely to play their dominant strategies in English auctions than in second-price auctions.

rium outcomes remain equilibrium outcomes in a first-price double auction when agents have full information. In this way, our paper shows that Bayesian mechanism design can be implemented, asymptotically, in a way that is robust in a number of dimensions that are relevant in practice.

Specifically, we study a two-sided private values setup with buyers who have unit demands and sellers who have unit capacities. We impose the assumption that it is sometimes but not always optimal for the market maker to induce all agents to trade. This assumption is innocuous insofar as without it the design problem is trivial. We assume that the market maker's objective is to achieve asymptotically the Bayesian objective of maximizing a weighted average of expected revenue and expected social surplus.[3] When the Bayesian optimal mechanism design problem has a unique local maximum, the Bayesian optimal mechanism induces trade by those agents who belong to the Walrasian set defined with respect to the weighted virtual types, which depend on the distributions. We assume that neither the market maker nor the agents know these distributions.

The key observation that we make and exploit is the connection between the empirical spacings of order statistics and the theoretical construct of virtual types. For buyers who draw their values independently from some distribution, we show that the empirical counterpart to the virtual value of the buyer with the kth highest draw is given by his value minus k times the spacing between the kth highest and the $k+1$st highest value. Analogously, the empirical counterpart to the virtual cost of the seller with the kth lowest draw is given by her cost plus k times the spacing between her cost and the next lowest cost.

The basic idea underlying our approach is then simple. The prices at which agents drop out in a double-clock auction (or, equivalently, the bids that they submit in a second-price auction) are statistically informative about the distributions from which types are drawn. This information can be used to estimate nonparametrically the virtual types needed to approximate the Bayesian optimal mechanism. Under the additional assumption that the Bayesian design problem is characterized by a unique local maximum (which is the case, for example, under Myerson's regularity condition), all agents who are inframarginal should trade once the virtual value is above the virtual cost. This means that incentive compatibility is not too difficult to handle even if, as is inevitable in a prior-free setup, virtual type functions have to be estimated, provided the estimates only depend on data

[3] A mechanism's ability to accommodate these two goals is valuable for market designers because their objectives often entail both. For example, the pivotal argument that led the U.S. Congress to allocate spectrum licenses using auctions was the U.S. Government's need for funds in the early 1990s, while the economists that paved the way for auctions to be used had emphasized the social surplus that such an allocation mechanism would generate. Similar arguments are being echoed in current discussions relating to the "incentive auction"; see, for example, Loertscher, Marx, and Wilkening (forthcoming) for more details.

inferred from agents who have dropped out.

At first sight, it might seem that one is done: Collect bids from agents who drop out and estimate distributions basically by counting. From these, infer the empirical virtual types of agents who have dropped out. Continue increasing the buyers' clock and decreasing the sellers' clock until the first point at which the empirical virtual value and virtual cost functions intersect. However, the problem with this intuitive idea is that the empirical virtual type functions are volatile, and the more so the less efficient is an agent's type. Consequently, in large economies the first point of intersection between the empirical virtual type functions is likely to occur close to the efficient quantity rather than at the Bayesian optimal one. This is so despite the fact that distributions are estimated with a high degree of accuracy as the number of agents becomes large. Intuitively, virtual type functions are volatile because they depend not only on the cumulative distribution, which will not vary much given a large number of observations, but also on the density, whose straightforward estimate remains volatile. We show that this problem can be overcome by appropriately smoothing the estimated virtual type functions using moving averages of spacings between order statistics. As the number of agents becomes large, the number of observations used in the estimation becomes large as well, but by letting the fraction of neighboring spacings included go to zero, the density and hence the virtual type function is accurately estimated at every point. Thereby, excessive volatility is eliminated, and the mechanism converges to the Bayesian optimum as the market becomes large.

The prior-free double-clock auction can be extended to accommodate heterogeneous groups of buyers and sellers, assuming that agents within each group draw their types from the same distribution, by running a separate clock for each group. In this extended, multi-clock auction, buyers and sellers in different groups face different prices, but those prices are synchronized so that virtual valuations across all buyer groups and virtual costs across all seller groups are the same. The multi-clock auction endows agents with obviously dominant strategies, is privacy preserving, and scales with market size just like the double-clock auction. It also preserves the property of envy-freeness, with the qualification that, because of price discrimination, this requirement is confined to agents within the same group.

Developing flexible two-sided exchanges that permit the reallocation of assets without having to renege on contractual promises associated with them is desirable and bound to become increasingly important because the accelerated pace of technological change increases uncertainty about the best use of many valuable resources, including, for example, radio spectrum and water-use rights. Trading on such exchanges will often be one-off events, which means that neither the designer nor the participants can rely upon past experiences or observations to gauge supply and demand or to formulate Bayesian optimal bidding strategies. This makes desirable the design of mechanisms that dispense with Bayesian notions, both for the rules of trade and for the equilibrium strategies. Consequently, for such novel environments and one-off reallocation problems for high-value assets, mechanisms with dominant strategies will be important and valuable, and the more so if the dominance property is transparent even to bidders with little or no prior experience. This is true for double-clock and multi-clock auctions because they endow agents with obviously dominant strategies in the sense defined by Li (2015).

Clock auctions also have the property of preserving the privacy of trading agents. This is important in order to avoid ex post regret for the designer and, in many instances with large sums at stake, for bidders to be willing to play their dominant strategies. For example, absent privacy preservation, bidders may fear that the information provided to the mechanism will be leaked to competitors or tax authorities, which may harm them further down the track.[4] As Milgrom and Segal (2015) observe, privacy preserving mechanisms have the additional benefit that trading agents need not have exact estimates of their valuations and costs, but need only know whether these are below or above some thresholds.

The envy-freeness of the equilibrium outcomes of our double-clock and multi-clock auctions is an important property in light of ubiquitous requirements for fair and nonarbitrary treatment of equals, in particular when the designer is Government.[5] Even when discrimination between certain groups of agents is permitted, designs may be subject to challenge based on being "arbitrary and capricious."[6]

Our double-clock auction allows a designer to delegate the setting of reserve prices, which are important for generating revenue and reducing procurement costs, to the mech-

[4]McMillan (1994) provides an early account of an auction design that did not preserve the privacy of bidders, which created political and public outcry about the money that was left on the table. See Ausubel (2004), Ausubel and Milgrom (2006), Rothkopf (2007), and Brandt and Sandholm (2008) for additional arguments for why privacy preservation matters.

[5]Regarding the provision of telecommunications services to schools and libraries, the U.S. Code of Federal Regulations states: "All potential bidders must have access to the same information and must be treated in the same manner." (Title 47 Telecommunication Section 54.642(b)(3)) According to Florida state law, the purpose of the bidding process, as asserted in the early landmark case of Wester v. Belote, 103 Fla. 976, 138 So. 721, 723-24 (1931), includes "to secure fair competition upon equal terms to all bidders" and "to afford an equal advantage to all desiring to do business with the county, by affording an opportunity for an exact comparison of bids."

[6]Although it is accepted that the U.S. Federal Communications Commission's spectrum license auctions may discriminate between large, small, and very small bidders through a program of bid credits for the small and very small bidders, the auctions have been challenged in U.S. courts in some instances based on being "arbitrary and capricious." (See "NAB Sues FCC Over Incentive Auction: Says FCC Auction Methodology Move was Illegal and Abuse of Discretion," Broadcasting & Cable, August 18, 2014, available at http://www.broadcastingcable.com/nab-sues-fcc-over-incentive-auction/133269, accessed February 24, 2015; and Rohde (2002, Chapter 1) regarding challenges raised to the FCC's approach to promoting minority ownership of spectrum.)

anism, which determines them as a function of the bid data. This addresses the challenge of how to set reserves when evidence regarding the relevant distributions is lacking—a challenge that is accentuated for government agencies that face requirements of nonarbitrary treatment and rules because any choice of distribution may be perceived as arbitrary. Moreover, no additional commitment by the designer is required provided he can commit to a double-clock auction and to the specific estimation that will be used by the mechanism and will determine his expectations. In particular, we show that for any weight on revenue in the designer's objective, the estimator can be chosen in such a way that it is self-enforcing for the designer to implement the optimal allocation rule when his expectations are determined by this estimator. That is, no commitment to the allocation rule is required.

Our double-clock auction and its multi-clock extension are flexible designs. They accommodate alternative objective functions, such as a revenue threshold beyond which the designer only cares about social surplus, and constraints such as caps on the number of units a certain subset of bidders can acquire or sell. Both revenue thresholds and restrictions on participation have been issues in the design of the U.S. Federal Communications Commission's incentive auction.[7] Our double-clock and multi-clock auctions also permit real-time diagnostics that inform the designer at any given point in time about whether revenue is more likely to increase or decrease if the procedure continues. Such feedback is essential when the designer faces hard ex post revenue thresholds because the feasibility of such thresholds is uncertain.

To address the question to what extent or in what sense our double-clock and multi-clock auctions are "optimal," we impose the constraint of privacy preservation of trading agents. We distinguish between what a Bayesian mechanism designer can achieve who faces the constraint of privacy preservation (and the usual incentive and individual rationality constraints), but can use his prior information about distributions, and what a mechanism designer can achieve who faces the same constraint, but has not the luxury of being endowed with a prior. This comparison is insightful for two reasons. First, as argued, privacy preservation is an important consideration in many real-world, large-scale allocation problems both for the participants and the designer. Second, as we show, privacy preservation is an operational constraint that can be imposed on allocation rules and mechanisms, which permits meaningful comparisons between different mechanisms. This is valuable because evaluating alternative designs in prior-free environments is inherently

[7]The incentive auction is required to raise revenue from buyers sufficient to cover payments to sellers plus certain other expenses (Public Law 112-96, Section 6403(c)(2)(B), http://www.gpo.gov/fdsys/pkg/PLAW-112publ96/pdf/PLAW-112publ96.pdf). Restrictions on participation by certain buyers has been raised by, among others, the U.S. Department of Justice, "Ex Parte Submission of the United States Department of Justice," WT Docket No. 12-269, April 11, 2013.

difficult.

With this in mind, we show that implementation via a double-clock auction is possible if and only if there is a corresponding direct mechanism that endows agents with dominant strategies, is envy free, and exhibits non-bossiness by trading agents, where non-bossiness means that trading agents affect neither the payments nor the allocations of other trading agents, conditional on trading. Because a direct mechanism that implements the efficient quantity, and more generally any Bayesian optimal quantity, violates non-bossiness, it follows immediately that no mechanism can implement the efficient or optimal quantity without violating privacy preservation. This impossibility result raises the question of what is the best one can do in the Bayesian setup under the constraint of privacy preservation. We answer this question by deriving the Bayesian optimal privacy preserving mechanism and by showing that an augmented version of our double-clock auction converges to the Bayesian optimum privacy preserving mechanism as estimation errors vanish.

This paper is at the intersection of detail-free, robust, and Bayesian mechanism design. It shares much of its motivation with Wilson (1987) and the subsequent literature on robust mechanism design in the tradition of Bergemann and Morris (2005, 2009, 2012).[8] It combines this with insights and elements from Bayesian mechanism design. Most notably, it relates to the contributions on optimal mechanisms for two-sided exchanges by Myerson and Satterthwaite (1983), Gresik and Satterthwaite (1989), and Williams (1999), which build on Myerson (1981).

Within the literature on dominant strategy mechanisms, the paper owes much to the contribution of McAfee (1992) and draws inspiration from Milgrom and Segal (2015). We first generalize McAfee's mechanism to accommodate Bayesian notions in the objective function, then we adapt it to a prior-free environment and relate it to Bayesian optimal privacy preserving mechanisms.[9] In our two-sided setup, we explore and exploit many of the same properties of double-cock auctions—such as privacy preservation, obviously dominant strategies, and weak-group strategy proofness—that Milgrom and Segal (2015) establish for clock auctions in one-sided environments. Although many of the distribution-free properties in the two papers are the same, their scopes and purposes are different and complementary. Milgrom and Segal develop one-sided deferred-acceptance auctions that can handle a multitude of computationally challenging technological constraints, whereas we develop a prior-free mechanism that is asymptotically Bayesian optimal and permits,

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

[9] Loertscher and Mezzetti (2015) provide a generalization of McAfee (1992) in a different direction by allowing for buyers and sellers whose types are multi-dimensional because they have multi-unit demands and supplies.

among other things, real-time feedback as to whether revenue is likely to increase if the double-clock auction continues.

Our paper also contributes to the literature on mechanism design with estimation.[10] The two most important precursors to the current paper within this strand of literature are Segal (2003) and Baliga and Vohra (2003). 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) define 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.[11]

There has also been an upsurge of interest in the computer science literature in dominant-strategy prior-free mechanisms along the lines of Baliga and Vohra (2003).[12] These mechanisms often have the desirable properties of convergence to the Bayesian optimal revenue and good worst-case performance relative to the benchmark of prior-free mechanisms that approximate Bayesian optimality but are not incentive compatible.[13] For example, Goldberg, Hartline, and Wright (2001) focus on dominant-strategy one-sided auctions for a good with unlimited supply. They define a mechanism that is similar to Baliga and Vohra's, study its worst-case performance, and show that it is within a constant factor of optimal single-price revenue.[14] Working with random sampling mechanisms, Dhangwatnotai, Roughgarden, and Yan (2015) show that for one-sided setups, even a randomly selected reserve price—set equal to the bid submitted by a randomly selected bidder—achieves at least half of the optimal revenue.

The 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

[10] There is a vast literature on estimation in auctions, which is notionally also related. We refer the reader to Athey and Haile (2007) and Guerre, Perrigne, and Vuong (2000) and the references therein. In Appendix B.1, we show that a natural alternative to our estimation approach—a kernel estimator à la Guerre, Perrigne, and Vuong (2000)—does not outperform our approach. In the small, it suffers from excessive volatility.

[11] As described in Baliga and Vohra (2003, Remark 2), it is not sufficient merely to estimate each agent's virtual valuation function by dropping only the agent's own report because then incentive compatibility is violated.

[12] Baliga and Vohra (2003) refer to their mechanism, which uses reports from a sample of agents to infer the distribution of types for other agents, as an "adaptive mechanism," whereas Goldberg, Hartline, and Wright (2001) and Goldberg et al. (2006), for example, refer to similar mechanisms as "random sampling mechanisms."

[13] See, e.g., Devanur, Hartline, and Yan (2015) on the formal definition of "approximate" in this sense.

[14] For additional work on benchmarks for evaluating the worst-case performance of two-sided mechanisms, see Deshmukh, Goldberg, Hartline, and Karlin (2002) and Dütting, Roughgarden, and Talgam-Cohen (2014).

(1973). How can a market maker infer the data necessary to clear the market without violating agents' incentive and participation constraints at no cost to himself? The short answer is that he cannot. However, one way of interpreting the results in McAfee (1992), Rustichini, Satterthwaite, and Williams (1994), Satterthwaite, Williams, and Zachariadis (2015) is that there are practical mechanisms that approximate full efficiency quickly as the economy grows. The problem that we address of maximizing revenue, or any convex combination of revenue and social surplus, is more complicated because the Bayesian optimal quantity is not a distribution-free concept, whereas the efficient quantity is.

The key feature that sets our paper apart from the previous and concurrent literature is that our mechanism is prior free and envy free for any market size and permits an implementation via clock auctions, which have, as argued, a number of desirable properties. Our methodological contribution is that we observe and exploit the connection between spacings of order statistics, which is an empirical measure, and virtual types, which is a theoretical concept. The direct empirical counterpart to virtual types depends only on one spacing and is thus too volatile for practical purposes. However, we show how excessive volatility can be overcome through appropriate averaging of spacings.

The remainder of this paper is organized as follows. In Section 2, we describe our setup and define the baseline prior-free mechanism. Section 3 shows the asymptotic optimality of the baseline prior-free mechanism. In Section 4, we turn to double-clock auctions, privacy preservation, the augmented prior-free mechanism, and performance in the small. In Section 5, we generalize the setup to allow for ex ante asymmetries among groups of buyers and groups of sellers. Section 6 contains extensions, and Section 7 concludes. The proofs are contained in Appendix A.

2 Model

In this section, we introduce the basic setup, review Bayesian mechanism design results that are relevant for our analysis, and introduce the baseline prior-free mechanism.

2.1 Setup

We consider a setup with a set of buyers $\mathbb{N}$ with unit demands and a set of sellers $\mathbb{M}$ with unit capacities. Throughout the paper, we denote the cardinality of a set $\mathbb{X}$ by lower-case x. We assume $n \geq 2$ and $m \geq 2$. Each buyer draws his value v independently from the distribution F, with support $[\underline{v}, \overline{v}]$ and positive density f. Each seller draws her cost c

from the distribution G, with support $[\underline{c}, \overline{c}]$ and positive density g. Let

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

denote the virtual value and virtual cost functions. We denote the buyers' values ranked in descending order by $v_{(1)} \geq ... \geq v_{(n)}$, and we denote the sellers' costs ranked in ascending order by $c_{[1]} \leq ... \leq c_{[m]}$, where ties are broken through randomization. Define $v_{(0)} \equiv \infty$, $v_{(n+1)} \equiv -\infty$, $c_{[0]} \equiv -\infty$, and $c_{[m+1]} \equiv \infty$.

Each agent is privately informed about his type, but the types and distributions from which they are drawn are unknown to the mechanism designer and the agents. The designer only knows that buyers and sellers draw their types independently from the same distributions, and that the distributions are such that the Bayesian design problem is characterized by a unique local maximum. In Section 5, we generalize the setup to the case in which there are different groups of buyers and sellers, with each group drawing from a different distribution. For the clock implementation, we assume that the designer knows a lower bound for $\underline{v}$ and an upper bound for $\overline{c}$, which allows the designer to start the clocks at prices that guarantee that all agents are active irrespective of their types.

Buyers, sellers, and the designer are risk neutral. A buyer's payoff is zero if he does not trade and is equal to his value minus the price he pays if he does trade. Similarly, a seller's payoff is zero if she does not trade and is equal to the payment she receives minus her cost if she does trade.

We assume that the designer wants to maximize a weighted sum of expected revenue and social surplus, with weight $\alpha \in [0, 1]$ on revenue, subject to agents' incentive compatibility and ex post individual rationality constraints. The case with $\alpha = 0$ corresponds to the objective of ex post efficiency, and the case with $\alpha = 1$ corresponds to the objective of revenue maximization. The weighted virtual types $\Phi_\alpha(v)$ and $\Gamma_\alpha(c)$ are defined as

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

By construction, when $\alpha = 0$, the weighted virtual types correspond to true types, and when $\alpha = 1$, to the well-known concept of a buyer's virtual valuation and the somewhat less familiar concept of a seller's virtual cost. 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 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, Marx, and Wilkening (forthcoming)).

2.2 Bayesian optimal mechanism benchmark

When distributions are known to the mechanism designer and the Bayesian design problem has a unique local maximum, the designer's objective is maximized by ranking buyers in increasing order by their weighted virtual values and sellers in decreasing order by their weighted virtual costs and having a pair trade if and only if the weighted virtual value exceeds the weighted virtual cost for that pair (see, e.g., Gresik and Satterthwaite, 1989). We refer to the mechanism that maximizes the designer's objective, conditional on incentive compatibility and ex post individual rationality, as the **α-optimal mechanism**.

We say that a direct mechanism satisfies **dominant strategies** if truthful reporting is a best response for each agent for all possible reports of the other agents. In the dominant strategy implementation of the α-optimal mechanism, each trading agent pays/receives the worst report he could have made without affecting the allocation. Thus, when there are q trades and virtual types are strictly increasing, trading buyers pay $\max\left\{v_{(q+1)}, \Phi_\alpha^{-1}\left(\Gamma(c_{[q]})\right)\right\}$ and trading sellers receive $\min\left\{c_{[q+1]}, \Gamma_\alpha^{-1}\left(\Phi_\alpha(v_{(q)})\right)\right\}$.

2.3 Baseline prior-free mechanism

As mentioned, the key assumption in our setup is that the distributions are not known to the mechanism designer or the agents. We now define a baseline prior-free mechanism that approximates the performance of the α-optimal mechanism and satisfies other desired properties.

Empirical virtual types and spacings between order statistics

Let $\hat{F}$ be the empirical distribution of values as revealed by the buyers' bids, with $\hat{F}(j) \equiv \frac{n+1-j}{n+1}$. Correspondingly, define $\hat{G}(j) \equiv \frac{j}{m+1}$. The empirical weighted virtual value and virtual cost for the jth highest value and jth lowest cost are

$$\hat{\Phi}_\alpha(j) \equiv v_{(j)} - \alpha \frac{1-\hat{F}(j)}{\frac{\hat{F}(j)-\hat{F}(j+1)}{v_{(j)}-v_{(j+1)}}} \quad \text{and} \quad \hat{\Gamma}_\alpha(j) \equiv c_{[j]} + \alpha \frac{\hat{G}(j)}{\frac{\hat{G}(j+1)-\hat{G}(j)}{c_{[j+1]}-c_{[j]}}},$$

where we approximate $f(v_{(j)})$ with the left-side derivative and $g(c_{[j]})$ with the right-side derivative based on the respective empirical distribution. Rewriting these expressions, we see that the virtual type functions of the jth traders are related to spacings between the

jth and $j+1$st order statistics as follows:

$$\hat{\Phi}_\alpha(j) = v_{(j)} - \alpha j \left(v_{(j)} - v_{(j+1)}\right) \quad \text{and} \quad \hat{\Gamma}_\alpha(j) = c_{[j]} + \alpha j \left(c_{[j+1]} - c_{[j]}\right). \tag{2}$$

The relationships in (2) connect the empirical concepts of order statistics with the theoretical construct of virtual type functions, which have been at the heart of Bayesian mechanism design ever since Myerson's path-breaking contribution.

Incentives

Let k be the efficient quantity and consider the following direct mechanism: trade the quantity $k-1$ and require buyers who trade to pay $v_{(k)}$ and pay $c_{[k]}$ to sellers who trade. An important insight from McAfee (1992) is that this mechanism endows agents with dominant strategies and respects their individual rationality constraints ex post. Extending this idea to include revenue in the designer's objective suggests using a similar mechanism, but where k is the highest indexed pair $(v_{(k)}, c_{[k]})$ such that $\hat{\Phi}_\alpha(k) \geq \hat{\Gamma}_\alpha(k)$. However, for any $\alpha > 0$ this simple idea suffers from the feature that for large numbers of buyers and sellers volatility in the empirical weighted virtual types is excessive. Because the number of traders is identified by the *largest* index such that the empirical weighted virtual value is greater than or equal to the empirical weighted virtual cost, with high probability this mechanism allows significantly more trades than would be optimal.

Figure 1(a) illustrates this for the case of $\alpha = 1$ (revenue maximization) and $n = m = 1000$ agents on each side whose types are drawn from the Uniform distribution on $[0, 1]$. As shown in the figure, the theoretical virtual types—which are derived from the theoretical distributions F and G using (1)—indicate an optimal quantity of 250 trades. However, the intersection of the empirical virtual types occurs at approximately twice as large a quantity. Indeed, the volatility of the empirical virtual types pulls the quantity traded based on the comparison of $\hat{\Phi}_\alpha(k)$ and $\hat{\Gamma}_\alpha(k)$ towards the efficient quantity rather than the α-optimal quantity.

As illustrated by Figure 1(a), the empirical weighted virtual types are correct on average, but problematic to use because of their volatility. Based on this observation, we define the baseline prior-free mechanism using the average of nearby spacings instead of the single spacing that appears in the empirical weighted virtual types. In addition, as described below, we adjust the coefficient multiplying the spacing so that the mechanism based on the smoothed virtual types is self-enforcing and deficit free.

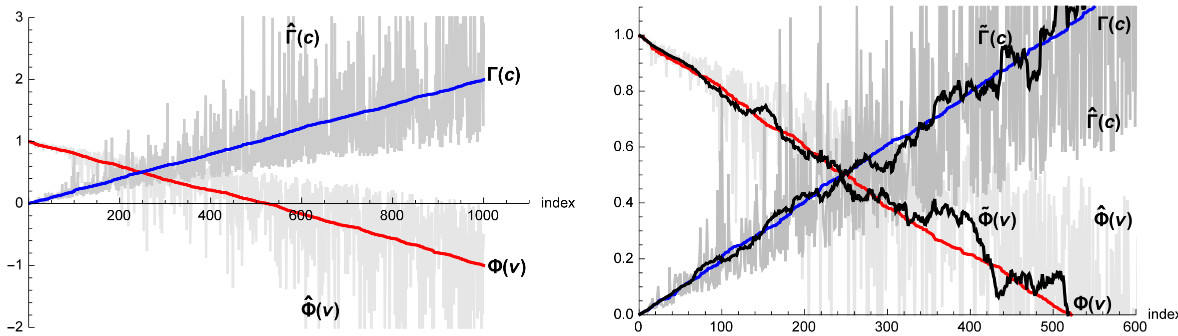

(a) Empirical virtual types

(b) Smoothed virtual types

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

Definition of the baseline prior-free mechanism

For $j \in \{1,...,n\}$ we define the jth smoothed weighted virtual value and for $j \in \{1,...,m\}$ we define the jth smoothed weighted virtual cost as follows:

$$\tilde{\Phi}_\alpha(j) \equiv v_{(j)} - \chi_\alpha(j)\sigma_j^v \quad \text{and} \quad \tilde{\Gamma}_\alpha(j) \equiv c_{[j]} + \chi_\alpha(j)\sigma_j^c,$$

where the coefficient $\chi_\alpha(j)$ is

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

and σ_j^v and σ_j^c are estimates of the spacing between values at $v_{(j)}$ and the spacing between costs at $c_{[j]}$, respectively, based on the average of nearby spacings:

$$\sigma_j^v \equiv \begin{cases} \frac{v_{(j)} - v_{(j+\min\{r_n, n-j\})}}{\min\{r_n, n-j\}}, & \text{if } j < n \\ \frac{1}{n+1}, & \text{otherwise} \end{cases} \quad \text{and } \sigma_j^c \equiv \begin{cases} \frac{c_{[j+\min\{r_m, m-j\}]} - c_{[j]}}{\min\{r_m, m-j\}}, & \text{if } j < m \\ \frac{1}{m+1}, & \text{otherwise,} \end{cases} \tag{3}$$

where r_x is $\sqrt{x}$ rounded to the nearest integer.[15] As explained in Section 2.4 below, the construction of $\chi_\alpha(j)$ is such that the mechanism is self-enforcing and deficit free. The choice of r_x such that $\lim_{x\to\infty} r_x = \infty$ and $\lim_{x\to\infty} \frac{r_x}{x} = 0$, along with continuity assumptions, ensures that σ_j^v converges in probability to $E[v_{(j)} - v_{(j+1)}]$ as n grows large and σ_j^c to $E[c_{[j+1]} - c_{[j]}]$ as m grows large. Combining these properties with the facts that $E\left[\left(1 - F(v_{(j)})\right)/f(v_{(j)})\right] = jE[v_{(j)} - v_{(j+1)}]$ and $E\left[G(c_{[j]})/g(c_{[j]})\right] = jE[c_{[j+1]} - c_{[j]}]$,

[15] Alternatively, r_x could be defined such that for some finite $\bar{x}$ and $\beta, \beta' \in (0,1)$ with $\beta < \beta'$, for all $x \geq \bar{x}$, $r_x \in (x^\beta, x^{\beta'})$.

and that the coefficient $\chi_\alpha(j)$ is of order j for $\alpha > 0$, gives us asymptotic optimality (see Section 3).

The **baseline prior-free mechanism** is defined as a direct mechanism in which all agents report their types. Its allocation and pricing rules are defined algorithmically as follows. For indices $t \in \{0, 1, ...\}$, define the state to be $\boldsymbol{\omega}_t = (e_t, k_t)$, with $e_t \in \{0, 1\}$ and $k_t \in \{0, 1, ...\}$. Initialize the algorithm with $e_0 \equiv 0$ and $k_t = \min\{n, m\}$. The allocation and pricing rules are determined by the lowest index $\tilde{t}$ such that $e_{\tilde{t}} = 1$, with the $k_{\tilde{t}}-1$ highest valuing buyers and lowest cost sellers trading at prices $v_{(k_{\tilde{t}})}$ and $c_{[k_{\tilde{t}}]}$. The transition from state $\boldsymbol{\omega}_t$ to $\boldsymbol{\omega}_{t+1}$ is defined by comparing the smoothed weighted virtual types. If $k_t = 0$ or $\tilde{\Phi}_\alpha(k_t) \geq \tilde{\Gamma}_\alpha(k_t)$, then $\boldsymbol{\omega}_{t+1} = (1, k_t)$ and otherwise $\boldsymbol{\omega}_{t+1} = (0, k_t - 1)$. Thus, the baseline prior-free mechanism trades $\tilde{k} - 1$ units, where $\tilde{k}$ is the largest index such that $\tilde{\Phi}_\alpha(\tilde{k}) \geq \tilde{\Gamma}_\alpha(\tilde{k})$, and the buyers who trade pay $v_{(\tilde{k})}$ and the sellers who trade are paid $c_{[\tilde{k}]}$.

As illustrated in Figure 1(b), the smoothed virtual types provide a good approximation to the theoretical virtual types in that example, and as shown below, the trading set defined by the smoothed virtual types approaches the trading set defined by the theoretical virtual types as the number of agents grows large. In Online Appendix B.1, we provide results using kernel-based weighted virtual types. The kernel-based mechanism considered there is also asymptotically optimal, but simulations suggest that it does not perform as well in the small as the baseline prior-free mechanism defined above.

2.4 Properties

The baseline prior-free mechanism is ex post individually rational and satisfies dominant strategies by the usual second-price logic. A trading agent's report does not affect the price he pays, and a nontrading agent cannot profitably deviate from truthful bidding because the resulting price would make trade unprofitable to any agent who would become a trader after a deviation when he would not trade under truthful bidding.

In addition, because an agent's bid does not affect the smoothed virtual types of lower indexed agents, the mechanism is **non-bossy** in the sense that no trading agent can, by changing his bid, affect the price or allocation of any other agent while still remaining an active trader.[16] Because all buyers pay the same price and all sellers pay the same price and those prices are determined by the value or cost of the most efficient nontrading agent, the mechanism is **envy free** in the sense that no agent prefers the allocation and price of another agent to his own. Because $\tilde{\Phi}_\alpha(j) \geq \tilde{\Gamma}_\alpha(j)$ implies $v_{(j)} \geq c_{[j]}$, the payments made by buyers exceed payments to sellers, so the mechanism is **deficit free** in the sense

[16] If a mechanism is non-bossy, then for generic types the TDI condition of Marx and Swinkels (1997) is satisfied for trading agents.

that the market maker never runs a deficit for any vector of types (or reports).

Like McAfee's mechanism, our baseline prior-free mechanism can be implemented via a **double-clock auction** that consists of an increasing buyers' clock and a decreasing sellers' clock and endows agents with dominant strategies. The starting prices of the two clocks are chosen so that initially all agents are active. Whenever the number of active buyers and sellers is not the same, only the clock on the long side of the market moves until the market is balanced. Agents' exit prices are used to determine smoothed virtual types $\tilde{\Phi}_\alpha$ and $\tilde{\Gamma}_\alpha$, and the double-clock auction stops when the number of remaining active agents on each side of the market is $\tilde{k}-1$, where $\tilde{k}$ is such that $\tilde{\Phi}_\alpha(\tilde{k}) \geq \tilde{\Gamma}_\alpha(\tilde{k})$ and $\tilde{\Phi}_\alpha(\tilde{k}+1) < \tilde{\Gamma}_\alpha(\tilde{k}+1)$. The implementation via a double-clock auction has a number of advantages over the direct mechanism as discussed in detail in Section 4, where we also provide a more formal definition of a double-clock auction. We only highlight two benefits of clock implementation here.

As is well-known, second-price auctions suffer from the problem that the bidders may not trust the designer that the prices they are charged or offered are truly the most competitive losing bids.[17] Provided a designer can commit to publicly observable clocks with monotone prices and activity rules, clock auctions and double-clock auctions provide a means of solving this credibility problem.

Moreover, if the designer can fix the methodology for estimating spacings based on the exit prices of agents who have dropped out, this may be all the commitment that is required because our prior-free double-clock auction is **self-enforcing** in the following sense.[18] Assume that the double-clock auction has proceeded to a point where $k-1$ agents remain active with the current clock prices being $v_{(k)}$ and $c_{[k]}$. To evaluate whether the market maker should stop the auction at these prices or continue to induce one additional exit on each side, he needs an estimate or predictor of the next exit prices. Given the spacings-based estimation procedure, the predicted prices are $\hat{v}_{(k-1)} \equiv v_{(k)} + \sigma_k^v$ and $\hat{c}_{[k-1]} \equiv c_{[k]} - \sigma_k^c$. The estimated change in the weighted objective, denoted $\Delta\hat{W}_{k-1}$, from continuing is thus the weighted sum of, first, the estimated increase in revenue from the $k-2$ trades that continue to occur minus the revenue from the one trade that no longer occurs and, second, the estimated loss in social surplus from the one trade that no longer

[17] See Lucking-Reiley (2000) for a colorful example.

[18] While alternative estimation procedures exist that could be applied to our problem, sticking to a given procedure does not appear to require much commitment unless it can be shown to be inferior to an alternative procedure. It is one thing to "commit" to something that has not been shown to be inferior and something else to commit to a policy that is obviously inferior ex post, such as charging five dollars instead of \$9.99 in a second-price auction when the second-highest bid was five and the highest was ten dollars.

occurs:

$$\Delta\hat{W}_{k-1} = \alpha\left(\left(k-2\right)\left(\sigma_k^v + \sigma_k^c\right) - \left(v_{(k)} - c_{[k]}\right)\right) - (1-\alpha)(v_{(k)} + \sigma_k^v - c_{[k]} + \sigma_k^c).$$

For $k \geq \frac{1+\alpha}{\alpha}$, we have $\Delta\hat{W}_{k-1} = -\left(\tilde{\Phi}_\alpha(k) - \tilde{\Gamma}_\alpha(k)\right)$, and for $k < \frac{1+\alpha}{\alpha}$, we have $\Delta\hat{W}_{k-1} = -\left(\tilde{\Phi}_\alpha(k) - \tilde{\Gamma}_\alpha(k)\right) + \left(\alpha(k-2) - (1-\alpha)\right)\left(\sigma_k^v + \sigma_k^c\right)$. Therefore, $\tilde{\Phi}_\alpha(k) \geq \tilde{\Gamma}_\alpha(k)$ implies $\Delta\hat{W}_{k-1} \leq 0$, and $\tilde{\Phi}_\alpha(k) < \tilde{\Gamma}_\alpha(k)$ implies $\Delta\hat{W}_{k-1} > 0$ for $k \geq \frac{1+\alpha}{\alpha}$ and $v_{(k)} < c_{[k]}$ otherwise. That is, if a designer has a preference parameter α and is constrained not to run a deficit, the allocation rule "trade $k-1$," where k the largest index such that $\tilde{\Phi}_\alpha(k) \geq \tilde{\Gamma}_\alpha(k)$, is self-enforcing (i.e. "sequentially rational" with expectations given by the mechanism).[19]

We summarize these and the previously observed properties in the following proposition.

Proposition 1 *The baseline prior-free mechanism is ex post individually rational, dominant strategy incentive compatible, non-bossy, envy free, and deficit free. It permits an implementation via a double-clock auction that is self-enforcing.*

Observe that the baseline prior-free mechanism is defined without reference to any distributional assumptions. All the properties summarized in Proposition 1 thus hold regardless of distributional properties.

3 Asymptotic optimality

We now show that the baseline prior-free mechanism is **asymptotically optimal** in the sense that, as the numbers of buyers and sellers go to infinity, the ratio of the value of the objective function under the mechanism to the value of the objective function under the α-optimal mechanism converges in probability to one. We now employ the assumption mentioned above that the virtual type functions defined in (1) are such that the Bayesian design problem is characterized by a unique local maximum. We also assume that the virtual type functions are continuous and that $\underline{v} \leq \overline{c}$. We prove the following result.

Proposition 2 *The baseline prior-free mechanism is asymptotically optimal.*

[19] The intuition for this self-enforcement is that the designer's expectations ("beliefs") are built into the mechanism and determined by the data and by the way spacings are estimated and weighted. With $\chi_\alpha(k)$ as the coefficient used in the smoothed virtual types, the objective is aligned with the predicted prices $\hat{v}_{(k-1)}$ and $\hat{c}_{[k-1]}$.

We illustrate the convergence result in Figure 2, which shows the ratio of baseline prior-free to optimal expected revenue for a variety of distributions and for n buyers and sellers, with n varying from 2 to 100.[20]

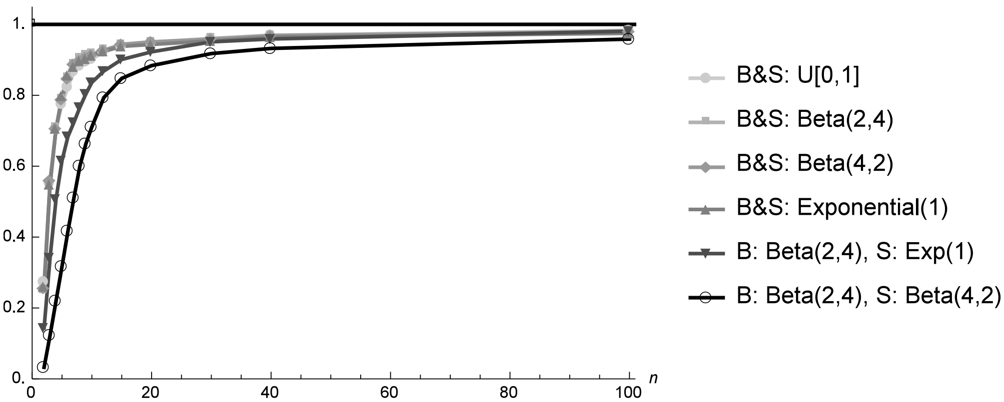

Figure 2: Ratio of baseline prior-free to optimal expected revenue for various numbers of symmetric agents based on Monte Carlo simulation (5000 auctions) with buyers' values and sellers' costs drawn from the distributions indicated.

To prove Proposition 2, we begin by showing in Lemma 1 that uniform bounds exist for the variance (denoted $V[\cdot]$) of $\chi_\alpha(j)\sigma_j^v$ and $\chi_\alpha(j)\sigma_j^c$ away from the boundary. Lemma 2 then shows that the difference between the theoretical and smoothed virtual values is uniformly convergent in probability to zero away from the boundary. Finally, Lemma 3 uses the preceding lemmas to show that the number of trades in the baseline prior-free mechanism approaches that in the optimal mechanism. Intuitively, if $\tilde{\Phi}_\alpha$ and $\tilde{\Gamma}_\alpha$ stay close to Φ_α and Γ_α, then the first intersection point of $\tilde{\Phi}_\alpha$ and $\tilde{\Gamma}_\alpha$ cannot be far from the intersection of Φ_α and Γ_α. Proposition 2 then follows from the fact that in both the optimal and the baseline prior-free mechanism it is the highest-valuing buyers and lowest-cost sellers that trade, and that in both cases the payments are in an interval bounded by the trading agent with the worst type and the nontrading agent with the best type and so differ by at most one spacing.

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

[20] To provide a sense of the error rate for our simulation results, for the case with $n = 5$ and values and costs drawn from the Uniform distribution on [0,1], simulations based on 5000 auctions generate a 95% confidence interval for the ratio of prior-free to optimal expected revenue that is plus or minus 1% of the mean.

Given Lemma 1, we can prove uniform convergence in probability of the theoretical and smoothed virtual types.

Lemma 2 *Given $\overline{\rho} \in (0,1)$, the difference between the theoretical and smoothed virtual values, $\Phi_\alpha(v_{(j)}) - \tilde{\Phi}_\alpha(j)$, on $[\overline{v} - \overline{\rho}(\overline{v} - \underline{v}), \overline{v}]$ is uniformly convergent in probability to zero. Similarly, on $[\underline{c}, \underline{c} + \overline{\rho}(\overline{c} - \underline{c})]$ the difference $\Gamma_\alpha(c_{[j]}) - \tilde{\Gamma}_\alpha(j)$ is uniformly convergent in probability to zero.*

Finally, given Lemma 2, we can now show that the number of trades in the baseline prior-free mechanism approaches that in the optimal mechanism.

Lemma 3 *As the number of agents goes to infinity, the share of agents that trade in the baseline prior-free mechanism converges in probability to the share in the optimal mechanism.*

For a discussion and evidence related to a rate of convergence on the order of $1/\min\{m,n\}$, see Online Appendix B.2.

4 Augmentation and Clock Implementation

That the baseline prior-free mechanism is asymptotically optimal and satisfies the desirable properties summarized in Proposition 1 is welcome news. However, it also raises the questions of how well the mechanism performs in the small and to what extent, or in what sense, the baseline prior-free mechanism or a variant thereof is optimal. In this section, we address both questions. We derive a benchmark, the Bayesian optimal privacy preserving mechanism, against which prior-free mechanisms can be evaluated, and we show that an augmented version of our mechanism achieves this benchmark except for the inevitable estimation errors. Then we show, using both simulations and analytical results, that this augmented prior-free mechanism performs well in the small. Because clock auctions preserve the privacy of trading agents, we begin by defining double-clock auctions formally. Then we relate double-clock auctions to non-bossiness and other properties of direct mechanisms. For the purposes of this section, we assume that virtual types are increasing and differentiable. The results can be generalized at the cost of additional complexity.

4.1 Preliminaries

To formally define a double-clock auction, we adapt the definition of a clock auction in Milgrom and Segal (2015) to our two-sided setting, in which one needs to ensure that the number of active buyers and sellers is the same at the time the procedure ends. In a double-clock auction, active agents may exit as the buyer clock price increases and the seller clock price decreases, but agents who exit remain inactive forever after exiting.

At period $t \in \{0, 1, ...\}$ of a **double-clock auction**, the state is $\boldsymbol{\omega}_t = (e_t, \boldsymbol{\omega}_t^B, \boldsymbol{\omega}_t^S)$ where $e_t \in \{0, 1\}$ specifies whether the double-clock auction has ended ($e_t = 1$) or not ($e_t = 0$), and $\boldsymbol{\omega}_t^B = (\mathbb{N}^A, \mathbf{x}^B, p^B)$ and $\boldsymbol{\omega}_t^S = (\mathbb{M}^A, \mathbf{x}^S, p^S)$ are buyer and seller states. The components of the buyer state are: the set of active buyers $\mathbb{N}^A \subseteq \mathbb{N}$ with cardinality n^A, the vector of exit prices for nonactive buyers $\mathbf{x}^B \in \mathcal{R}^{n-n^A}$, and the buyer clock price $p^B \in \mathcal{R}$. The seller state has an analogous structure. Let Ω be the set of all possible states. A double-clock auction starts in state $\boldsymbol{\omega}_0 \equiv (0, \boldsymbol{\omega}_0^B, \boldsymbol{\omega}_0^S)$, where $\boldsymbol{\omega}_0^B = (\mathbb{N}, \emptyset, \underline{p}^B)$ and $\boldsymbol{\omega}_0^S = (\mathbb{M}, \emptyset, \overline{p}^S)$ with $\underline{p}^B \leq \underline{v}$ and $\overline{p}^S \geq \overline{c}$.

A double-clock auction is defined by comparison functions $\phi : \mathcal{R} \times \Omega \to \mathcal{R}$ and $\gamma : \mathcal{R} \times \Omega \to \mathcal{R}$ and a target function $\tau : \Omega \to \mathcal{R}$. The comparison and target functions determine the transition from state $\boldsymbol{\omega}_t$ to $\boldsymbol{\omega}_{t+1}$. A double-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 buyers' clock price and the active sellers receiving the sellers' clock price.

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

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

If $n^A > m^A$, increase the buyer clock from p^B until a buyer i exits at $\hat{p}^B$. Set $\boldsymbol{\omega}_{t+1}^B = (\mathbb{N}^A \backslash i, (\mathbf{x}^B, \hat{p}^B), \hat{p}^B)$, $\boldsymbol{\omega}_{t+1}^S = \boldsymbol{\omega}_t^S$, and $e_{t+1} = 0$.

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

Properties of double-clock auctions

A double-clock auction is **privacy preserving** for agents who trade insofar as their types are never revealed to the mechanism (or any outside observer for that matter). As discussed in the introduction, privacy preservation is an important attribute in practice, particularly for large-scale, one-off allocation problems.

In our environment with unit demands and unit supplies, double-clock auctions also have **obviously dominant strategies** as defined by Li (2015): 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. Robustness with respect to the fine details of the environment makes it desirable for mechanisms to endow agents with dominant strategies. Because endowing agents with dominant strategies and having agents recognize their dominant strategies are two distinct things in practice, the obviousness of the dominant strategies in clock auctions is another powerful argument for their use and for focusing on direct mechanisms that can be implemented via clock auctions.

In addition, a double-clock auction is **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.[21] This property implies that, in the absence of transfers, collusion among a subset of agents cannot be strictly profitable for all of the colluding bidders.

Like the baseline prior-free mechanism, double-clock auctions are non-bossy. Non-bossiness is a restriction on direct and indirect mechanisms, whereas privacy preservation can only be satisfied by a subset of indirect mechanisms. The connection between the two properties is that an indirect mechanism is privacy preserving if and only if the equivalent direct mechanism satisfies non-bossiness, where "equivalent" means that for any given realization of types it makes the same transfers and induces the same allocation in equilibrium.

We summarize key properties of a double-clock auction and their relation to properties of direct mechanisms in the following proposition.

Proposition 3 *A direct mechanism satisfies dominant strategies, non-bossiness, and envy-freeness if and only if it can be implemented via a double-clock auction. Further, a double-clock auction is privacy preserving, endows agents with obviously dominant strategies, and is weakly group strategy-proof.*

The following proposition, which provides an impossibility result, follows from Proposition 3 and the fact that any α-optimal mechanism depends on the types of the marginal trading agents and so violates non-bossiness.[22]

[21] On weak group strategy-proofness in a one-sided clock auction, see Milgrom and Segal (2015). On the connection between individual and group strategy-proofness, see Barberà, Berga, and Moreno (2014).

[22] To see this, fix $(\boldsymbol{v}, \boldsymbol{c})$ and note that the α-optimal mechanism trades a quantity $q_\alpha(\boldsymbol{v}, \boldsymbol{c})$ such that

$$\Phi_\alpha(v_{(q_\alpha(\boldsymbol{v},\boldsymbol{c}))}) \geq \Gamma_\alpha(c_{[q_\alpha(\boldsymbol{v},\boldsymbol{c})]}) \quad \text{and} \quad \Phi_\alpha(v_{(q_\alpha(\boldsymbol{v},\boldsymbol{c})+1)}) < \Gamma_\alpha(c_{[q_\alpha(\boldsymbol{v},\boldsymbol{c})+1]}),$$

Proposition 4 *No α-optimal mechanism can be implemented while preserving the privacy of trading agents.*

Because α-optimal mechanisms violate non-bossiness, Propositions 3 and 4 imply that in the two-sided setting no α-optimal mechanism can be implemented as a double-clock auction. By setting $\alpha = 0$, Proposition 4 also establishes the impossibility of ex post efficient trade from a different and, to our knowledge, novel angle: If one insists on preserving the privacy of agents who trade, then it is impossible to always trade the efficient quantity in a two-sided setup provided only that full trade is sometimes but not always efficient. This contrasts with setups with private information only on one side of the market, where the English auction and its extension to multiple units by Ausubel (2004) permit ex post efficiency under privacy preservation. The conditions under which the impossibility result of Proposition 4 holds are remarkably general compared to the conditions for ex post efficiency to be possible or impossible for general two-sided trading environments (see Williams (1999) for a characterization for such environments). It only requires overlapping supports for the buyers' and the sellers' distributions, and it is not restricted to ex post efficiency.[23] It is also noteworthy that this impossibility has nothing to do with individual rationality constraints (be they ex ante, interim, or ex post), which are key for existing impossibility results.

Because impossibility results and non-achievable benchmarks feature prominently in the related computer science literature, it is useful to highlight the main differences between the two approaches and the results they generate. As mentioned, the focus of the computer science literature is on worst-case performance relative to non-incentive compatible benchmarks. Accordingly, this literature evaluates mechanisms based on their ability to perform within a constant factor of those benchmarks. For example, Deshmukh, Goldberg, Hartline, and Karlin (2002) compare revenue from prior-free, dominant strategy, two-sided mechanisms to the benchmark of the maximum revenue that could be achieved by an omniscient designer who knows the bidders' types and uses a single buyer price and a single seller price, subject to ex post individual rationality (but not incentive compatibility).[24] In contrast, we examine dominant-strategy incentive compatible mechanisms

which is unique almost surely when types are continuously distributed. Dropping the dependence of q_α on the profile of reports for notational ease, then given q_α, buyer i trades if and only if $v_i \geq \max\left\{v_{(q_\alpha+1)}, \Phi_\alpha^{-1}\left(\Gamma_\alpha(c_{[q_\alpha]})\right)\right\}$. But the right side depends on the report of the trading seller with type $c_{[q_\alpha]}$. Thus, the α-optimal quantity cannot be implemented without violating non-bossiness.

[23] Again, this contrasts with the one-sided setting. For example, with n ex ante symmetric buyers who draw their types independently from a known distribution satisfying Myerson's regularity condition and $k < n$ items for sale of zero value to the seller, the clock would be initialized at the price at which the weighted virtual value is zero and then increased until $n - k$ buyers have dropped out.

[24] Dütting, Roughgarden, and Talgam-Cohen (2014) examine the worst-case efficiency of two-sided mechanisms satisfying various properties, including dominant strategies, deficit freeness, and weak group

that preserve the privacy of agents who trade. As just argued, α-optimal mechanisms of this form exist for one-sided setups but not for two-sided environments.

To develop a benchmark for our mechanism, we next analyze what is the best a Bayesian mechanism designer can achieve subject to privacy preservation, and then compare the performance of a prior-free mechanism to this Bayesian optimal benchmark.

4.2 Optimal privacy preserving mechanisms

Given that any α-optimal mechanism violates privacy preservation, we now characterize the optimal mechanism for a designer with preference parameter α who is endowed with a correct prior F and G but, in addition to individual rationality and incentive constraints, has to respect the privacy of agents who trade.

Bayesian optimal privacy preserving mechanism

A good starting point is McAfee (1992), whose focus is on ex post efficiency. Letting k be the largest integer such that $v_{(k)} \geq c_{[k]}$, his mechanism trades k units at the price $p = \frac{v_{(k+1)}+c_{[k+1]}}{2}$ if $v_{(k)} \geq p \geq c_{[k]}$, and otherwise $k-1$ units at prices $p^B = v_{(k+1)}$ and $p^S = c_{[k+1]}$. This mechanism thus trades one more unit than the baseline prior-free mechanism in some circumstances. Importantly, it permits a double-clock auction implementation that respects the privacy of the agents who trade. This example shows that, for $\alpha = 0$, the baseline prior-free mechanism is not optimal in the domain of privacy preserving mechanisms.

In light of Proposition 4, it is natural to ask what is the best one can do subject to preserving the privacy of trading agents. To address this question, we temporarily stipulate that the designer knows F and G, but is restricted to preserve the privacy of trading agents, and we derive the mechanism that, conditional on $\Phi_\alpha(v_{(j)}) \geq \Gamma_\alpha(c_{[j]})$, maximizes the probability that the quantity traded is j given that the types $v_{(j+1)}, ..., v_{(n)}$ and $c_{[j+1]}, ..., c_{[m]}$ have been revealed and are precluded from trading and that no other types have been revealed. We refer to such a mechanism as the **Bayesian optimal privacy preserving (BOPP)** mechanism (which also satisfies dominant strategies).

Given $v_{(j+1)}, ..., v_{(n)}$, $c_{[j+1]}, ..., c_{[m]}$, and $\Phi_\alpha(v_{(j+1)}) < \Gamma_\alpha(c_{[j+1]})$, any BOPP mechanism must maximize the probability that, conditional on $\Phi_\alpha(v_{(j)}) \geq \Gamma_\alpha(c_{[j]})$, the quantity traded is j. This boils down to choosing, for all j, the target $\theta_j^* \in [\Phi_\alpha(v_{(j+1)}), \Gamma_\alpha(c_{[j+1]})]$ that maximizes the probability that $\Gamma_\alpha(c_{[j]}) \leq \theta_j^* \leq \Phi_\alpha(v_{(j)})$. A double-clock auction can determine whether $\Gamma_\alpha(c_{[j]}) \leq \theta_j^* \leq \Phi_\alpha(v_{(j)})$ in a privacy preserving way by stopping the

strategy-proofness, relative to the efficient outcome and provide bounds on the constant factor that can be achieved. For related work in the one-sided setup, see Chen, Gravin, and Lu (2014) and Devanur, Hartline, and Yan (2015).

procedure if the buyer clock reaches p^B such that $\Phi_\alpha(p^B) = \theta_j^*$ and the seller clock reaches p^S such that $\Gamma_\alpha(p^S) = \theta_j^*$ before the jth buyer and seller exit. As shown in the proof of Proposition 5, the optimal choice for θ_j^* is

$$\theta_j^* \equiv \begin{cases} \Gamma_\alpha(c_{[j+1]}), & \text{if } J(\theta) < 0 \text{ for all } \theta \in [\Phi_\alpha(v_{(j+1)}), \Gamma_\alpha(c_{[j+1]})] \\ \theta^*, & \text{if } J(\theta^*) = 0 \text{ for some } \theta^* \in [\Phi_\alpha(v_{(j+1)}), \Gamma_\alpha(c_{[j+1]})] \\ \Phi_\alpha(v_{(j+1)}), & \text{if } J(\theta) > 0 \text{ for all } \theta \in [\Phi_\alpha(v_{(j+1)}), \Gamma_\alpha(c_{[j+1]})], \end{cases} \tag{4}$$

where

$$J(\theta) \equiv \frac{1 - F(\Phi_\alpha^{-1}(\theta))}{f(\Phi_\alpha^{-1}(\theta))} \frac{1}{\Phi_\alpha^{-1\prime}(\theta)} - \frac{G(\Gamma_\alpha^{-1}(\theta))}{g(\Gamma_\alpha^{-1}(\theta))} \frac{1}{\Gamma_\alpha^{-1\prime}(\theta)}.$$

Any BOPP mechanism is then defined by comparison and target functions given by

$$\phi(p^B, \boldsymbol{\omega}_t) = \Phi_\alpha(p^B) \quad \text{and} \quad \gamma(p^S, \boldsymbol{\omega}_t) = \Gamma_\alpha(p^S) \quad \text{and} \quad \tau(\boldsymbol{\omega}_t) = \theta_{n^A}^*.$$

Proposition 5 *Any BOPP mechanism is characterized by a target function $\tau(\boldsymbol{\omega}_t) = \theta_{n^A}^*$, where $\theta_{n^A}^*$ is defined in* (4).

Proposition 5 has an intuitive explanation. The closest analog is the problem faced by a market maker who makes take-it-or-leave-it price offers p^B and p^S to a buyer and a seller—in our case, these are the n^A most efficient agents on each side—who draw their types from distributions H and L. Trade occurs if and only if both the buyer and the seller accept. Under quasiconcavity, the probability of trade is maximized when

$$\frac{L(p^S)}{l(p^S)} = \frac{1 - H(p^B)}{h(p^B)}.$$

This equates the semi-elasticity of supply l/L with the semi-elasticity of demand $h/(1 - H)$. In our setup, the prices p^B and p^S are required to satisfy $p^B = \Phi_\alpha^{-1}(\theta)$ and $p^S = \Gamma_\alpha^{-1}(\theta)$ with $\theta \in [\Phi_\alpha(v_{(n^A+1)}, \Gamma_\alpha(c_{[n^A+1]})]$, which explains the additional derivatives that multiply the semi-elasticities in the first-order condition defining θ^*. For example, for F uniform on $[\underline{v}, \overline{v}]$ and G uniform on $[\underline{c}, \overline{c}]$,[25] we have $\theta^* = \frac{\overline{v} - \underline{c}}{2}$, so $\tau(\boldsymbol{\omega}_t) = \min\left\{\Gamma_\alpha(c_{[n^A+1]}), \max\left\{\Phi_\alpha(v_{(n^A+1)}), \frac{\overline{v}-\underline{c}}{2}\right\}\right\}$.

Augmented prior-free mechanism

Guided by the BOPP mechanism, we now augment the baseline prior-free mechanism to allow an additional trade under certain circumstances. To define this **augmented prior-free mechanism**, it will be useful to identify the reliance of the jth smoothed weighted

[25]In this case, $\Phi_\alpha(x) = x(1 + \alpha) - \alpha\overline{v}$ and $\Gamma_\alpha(x) = (1 + \alpha)x - \alpha\underline{c}$.

virtual types on the jth value or cost separately from their reliance on the index j, so we will sometimes write the smoothed weighted virtual types as

$$\tilde{\Phi}_\alpha(j;p) \equiv p - \chi_\alpha(j)\sigma_j^v \text{ and } \tilde{\Gamma}_\alpha(j;p) \equiv p + \chi_\alpha(j)\sigma_j^c.$$

The augmented prior-free double-clock auction is defined by comparison and target functions given by

$$\phi(p^B, \boldsymbol{\omega}_t) = \tilde{\Phi}_\alpha(n^A+1;p^B) \quad \text{and} \quad \gamma(p^S, \boldsymbol{\omega}_t) = \tilde{\Gamma}_\alpha(m^A+1;p^S) \quad \text{and} \quad \tau(\boldsymbol{\omega}_t) = \tilde{\theta}^*_{n^A},$$

where $\tilde{\theta}^*_j \equiv 0$ for $j = \min\{n,m\}$ and otherwise

$$\tilde{\theta}^*_j \equiv \begin{cases} \tilde{\Gamma}_\alpha(j+1), & \text{if } \tilde{J}(\theta) < 0 \text{ for all } \theta \in [\tilde{\Phi}_\alpha(j+1), \tilde{\Gamma}_\alpha(j+1)] \\ \tilde{\theta}, & \text{if } \tilde{J}(\tilde{\theta}) = 0 \text{ for some } \tilde{\theta} \in [\tilde{\Phi}_\alpha(j+1), \tilde{\Gamma}_\alpha(j+1)] \\ \tilde{\Phi}_\alpha(j+1), & \text{if } \tilde{J}(\theta) > 0 \text{ for all } \theta \in [\tilde{\Phi}_\alpha(j+1), \tilde{\Gamma}_\alpha(j+1)], \end{cases}$$

where

$$\tilde{J}(\theta) \equiv \frac{\theta - \tilde{\Phi}_\alpha(j+1)}{(1+\alpha)\sigma_{j+1}^v} - \frac{\tilde{\Gamma}_\alpha(j+1) - \theta}{(1+\alpha)\sigma_{j+1}^c}.$$

Intuitively, $\tilde{\theta}^*_j$ equates the change in smoothed weighted virtual types required to get the virtual value up to $\tilde{\theta}^*_j$ and the virtual cost down to $\tilde{\theta}^*_j$, adjusted for the estimated change in virtual types between order statistics, $(1+\alpha)\sigma_{j+1}^v$ and $(1+\alpha)\sigma_{j+1}^c$.

As with the baseline prior-free mechanism, the allocation and prices depend only on reports of agents who do not trade. Thus, like the baseline prior-free mechanism, the augmented prior-free mechanism satisfies the properties of Proposition 1. As we elaborate in more detail below, it has the advantage of performing well in the small. As a corollary of Proposition 3, it has a privacy preserving, obviously dominant strategies, weak group strategy-proof implementation as a double-clock auction that is self-enforcing in the same way as the baseline prior-free mechanism.

Proposition 6 *The augmented prior-free mechanism is asymptotically optimal and satisfies the properties of the baseline prior-free mechanism stated in Proposition* 1.

Absent knowledge of the underlying distributions, it is not a priori clear how to compare alternative mechanisms that are prior free and satisfy all the other desirable criteria. Moreover, alternative mechanisms may vary in how they estimate distributions and virtual types. For example, we could use a moving average that assigns differential rather than uniform weights to the observations it uses. Accordingly, the performance of alternative mechanisms may vary with the underlying distributions because they use different estimation procedures.[26] In evaluating alternative designs, it is therefore useful to distinguish between how a mechanism performs with and without the estimation errors that are unavoidable in any prior-free mechanism. The augmented prior-free mechanism would achieve the outcomes of the BOPP mechanism were it not for estimation error.

4.3 Performance in the small

Although the augmentation provided by an optimally chosen target function does not affect the asymptotic properties of the mechanism, it can be significant in small markets. To see the impact of an optimally chosen target function in the small, define the fixed-target BOPP mechanism to be the same as the BOPP mechanism, but with the suboptimally chosen target function defined by $\theta_j \equiv \left(\Phi_\alpha\left(v_{(j+1)}\right) + \Gamma_\alpha\left(c_{[j+1]}\right)\right)/2$. This is a generalization of the mechanism proposed by McAfee (1992) to weighted virtual types. As shown in Proposition 7, the probability that an optimally chosen θ_j results in the efficient number of trades, whereas the fixed-target version comes up short by one, can be as large as 1/4 when the support of the value distribution extends sufficiently above that of the cost distribution. In the following proposition, we say that a distribution $H(x)$ with support $[0,1]$ is symmetric if $H(x) = 1 - H(1-x)$ holds for all $x \in [0,1]$.

Proposition 7 *Assume $\alpha = 0$. Let G and F be symmetric distributions on $[0,1]$. If costs are drawn from G and values from $F(\cdot/\overline{v})$ with support $[0,\overline{v}]$, then as $\overline{v}$ increases, the probability that the BOPP mechanism is efficient but the fixed-target BOPP mechanism is not goes to $\frac{1}{2^{\min\{n,m\}}}$.*

Figure 3(a) shows the performance of the baseline and augmented prior-free mechanism for $\alpha = 0$ relative to the optimal mechanism and the BOPP mechanism. The figure illustrates the performance enhancement achieved using the augmentation to execute an additional trade when possible while still maintaining privacy preservation. It also shows that the augmented prior-free mechanism achieves 87% or more of the BOPP outcomes, even when there are as few as 2 buyers and 2 sellers.

For the case of $\alpha = 1$, we have already seen in Figure 2 that for a range of distributions, average revenue of the baseline prior-free mechanism is at least 90% of optimal expected revenue when there are 30 or more buyer-seller pairs. To illustrate results for small markets with α away from the extremes, Figure 3(b) shows the augmented prior-free mechanism's performance relative to the α-optimal mechanism for various α. As shown in Figure 3(b), generally speaking, the smaller is α, the smaller is the impact of estimation error, and so the better is the performance of the augmented prior-free mechanism. However, the

[26] Other methods of estimation include kernel density estimation (see Online Appendix B.1).

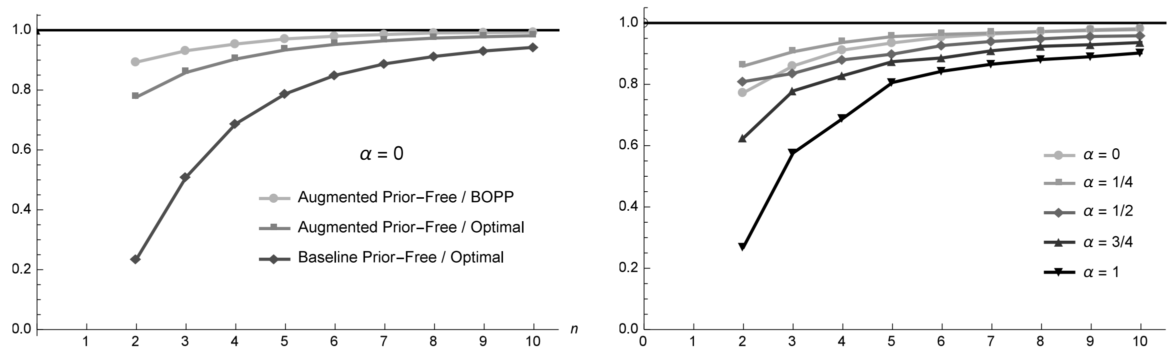

(a) Relative performance of mechanisms for $\alpha = 0$ (b) Effects of α on the baseline prior-free

Figure 3: Panel (a): Ratio of the prior-free to the optimal and BOPP outcomes for $\alpha = 0$. Panel (b): Ratio of the baseline prior-free to the α-optimal expected weighted objective for various weights α on revenue. Both panels assume $n = m$ and are based on Monte Carlo simulation (5000 auctions) with values and costs drawn from the Uniform distribution on $[0, 1]$.

augmented prior-free mechanism's use of the first buyer and seller exits for estimation is a greater disadvantage relative to the optimal mechanism when the number of agents is small. Thus, for small numbers of agents, the relative performance of the augmented prior-free mechanism can be better for larger α.

We provide a characterization of how the augmented prior-free mechanism relates to the BOPP mechanism for small markets in Proposition 8.

Proposition 8 *Assume $\alpha = 0$. Given any type realization, the number of trades in the augmented prior-free mechanism is within one of the number of trades in the BOPP mechanism, and if $\underline{v} < \overline{c}$, for a positive measure set of type realizations, the augmented prior-free mechanism produces greater social surplus than the BOPP mechanism.*

For very small numbers of agents ($n = m = 2$), Table 1 provides results for alternative distributions and mechanisms for $\alpha = 0$ by comparing outcomes under the baseline prior-free, augmented prior-free, BOPP, and optimal mechanisms. The higher values for the augmented prior-free mechanism relative to the baseline prior-free mechanism illustrate the value of augmentation.

Because the efficient quantity is a distribution-free concept, it seems natural to conjecture that estimation only matters for $\alpha > 0$. Interestingly, this is not the case. As we illustrate in Online Appendix B.3, estimation can improve outcomes because it allows the target function to vary with observable details, which is relevant even with $\alpha = 0$.[27]

[27] Intuitively, for $\alpha = 0$ the target function plays the role of the posted price in Hagerty and Rogerson (1987). With estimation, this posted price becomes a function of the distributions.

	U[0,1]	Exp[1]	Beta[2,2]	B:Beta[2,4] S:Beta[4,2]
BOPP / Optimal	87%	79%	77%	60%
Augmented Prior−Free / BOPP	89%	99%	99%	92%
Augmented Prior−Free / Optimal	77%	78%	76%	55%
Baseline Prior−Free / Optimal	26%	26%	25%	3%

Table 1: Performance in the small ($n = m = 2$): Comparison of baseline prior-free, augmented prior-free, BOPP, and optimal mechanisms with $\alpha = 0$ for different distributions.

5 Generalization to heterogeneous groups

In this section, we show that the augmented prior-free mechanism can be extended to accommodate known ex ante asymmetries among buyers and among sellers. As in the previous section, we assume that virtual types are increasing and differentiable, although the results can be generalized at the cost of additional complexity.

5.1 Setup

As before, we let $\mathbb{N}$ and $\mathbb{M}$ denote the sets of buyers and sellers with cardinalities n and m. We now allow, without requiring, the possibility that agents have characteristics that are observable to the designer, so that the designer can a priori place subsets of agents into groups of symmetric agents while allowing for asymmetries across different groups. For example, traders of carbon emission permits might be identifiable as either power plants, cement manufacturers, or other manufacturers, with traders within a group being symmetric, but with the possibility of asymmetries across groups. Let $\mathbb{Z}^B$ and $\mathbb{Z}^S$ be the sets of group labels for buyers and sellers. We refer to the setup in which $|\mathbb{Z}^B| = |\mathbb{Z}^S| = 1$ studied thus far as the **symmetric setup**. Let $n^b \geq 1$ be the number of buyers in buyer group b and let $m^s \geq 1$ be the number of sellers in group s, where $n = \sum_{b\in\mathbb{Z}^B} n^b$ and $m = \sum_{s\in\mathbb{Z}^S} m^s$. We assume that at least one buyer group and at least one seller group has 2 or more members. The group membership of each buyer and seller is common knowledge.

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

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

are increasing in their arguments for each buyer and seller group. For the dynamic implementation, we assume that the designer knows a lower bound for $\min_{b\in\mathbb{Z}^B}\{\underline{v}_b\}$ and an upper bound for $\max_{s\in\mathbb{Z}^S}\{\overline{c}_s\}$, which allows the designer to start the clocks at prices that guarantee that all agents are active irrespective of their types.

Group-specific weighted virtual types are given by

$$\Phi^b_\alpha(v) \equiv \alpha\Phi^b(v)+(1-\alpha)v = v-\alpha\frac{1 - F^b(v)}{f^b(v)} \quad \text{and} \quad \Gamma^s_\alpha(c) \equiv \alpha\Gamma^s(c)+(1-\alpha)c = c+\alpha\frac{G^s(c)}{g^s(c)}.$$

Under the stipulated assumptions, the Bayesian optimal mechanism induces trade by those agents who belong to the Walrasian set defined with respect to the weighted virtual types $\left(\Phi^b_\alpha(v_i)\right)_{i\in\mathbb{N}}$ and $\left(\Gamma^s_\alpha(c_j)\right)_{j\in\mathbb{M}}$. Group-specific smoothed weighted virtual types are given by $\tilde{\Phi}^b_\alpha(j) \equiv v^b_{(j)} - \chi_\alpha(j)\sigma^b_j$ and $\tilde{\Gamma}^s_\alpha(j) \equiv c^s_{[j]} + \chi_\alpha(j)\sigma^s_j$, where σ^b_j and σ^s_j are spacing estimates for each group defined analogously to the symmetric setup.

5.2 Mechanisms

We first extend the definition of a double-clock auction to a multi-clock auction, and then generalize the BOPP and augmented prior-free mechanisms.

Multi-clock auction

For the setup with heterogeneous groups, we define a multi-clock auction with separate price clocks for each buyer group and each seller group. Compared to the double-clock auction in the symmetric setup, the buyer and seller state in period t is extended to $\boldsymbol{\omega}^B_t = (\boldsymbol{\omega}^b_t)_{b\in\mathbb{Z}^B}$ and $\boldsymbol{\omega}^S_t = (\boldsymbol{\omega}^s_t)_{s\in\mathbb{Z}^S}$, where the group states $\boldsymbol{\omega}^b_t$ and $\boldsymbol{\omega}^s_t$ are defined analogously to before. With this adjusted notation at hand, the state in period t is still given as $\boldsymbol{\omega}_t = (e_i, \boldsymbol{\omega}^B_t, \boldsymbol{\omega}^S_t)$. The comparison functions ϕ^b and γ^s are defined for each group. For $t \in \{0, 1, ...\}$, if for $b \in \mathbb{Z}^B$, $\boldsymbol{\omega}^b_t = (\mathbb{N}^{A^b}, \mathbf{x}^b, p^b)$, for $s \in \mathbb{Z}^S$, $\boldsymbol{\omega}^s_t = (\mathbb{M}^{A^s}, \mathbf{x}^s, p^s)$, and $e_t = 0$, then $\boldsymbol{\omega}_{t+1}$ is determined as follows:

If $\sum_{b\in\mathbb{Z}^B} n^{A^b} = \sum_{s\in\mathbb{Z}^S} m^{A^s}$: If $\sum_{b\in\mathbb{Z}^B} n^{A^b} = 0$ or $\min_{b\in\mathbb{Z}^B} \phi^b(p^b, \boldsymbol{\omega}_t) \geq \max_{s\in\mathbb{Z}^S} \gamma^s(p^s, \boldsymbol{\omega}_t)$, then $\boldsymbol{\omega}_{t+1} = (1, \boldsymbol{\omega}^B_t, \boldsymbol{\omega}^S_t)$. Otherwise, proceed as follows (the order in which clocks on either side of the market are moved is again immaterial):

Increase the vector of buyer clock prices from $\boldsymbol{p}^B = (p^b)_{b\in\mathbb{Z}^B}$ by increasing the clocks for the smallest number of buyer groups possible so as to increase $\min_{b\in\mathbb{Z}^B} \phi^b(\cdot, \boldsymbol{\omega}_t)$

(or by increasing the clock of only one buyer group if $\min_{b\in\mathbb{Z}^B} \phi^b(\cdot,\boldsymbol{\omega}_t)$ is nonincreasing) until either there is an exit by a group $\hat{b}$ buyer i at clock price vector $\hat{\mathbf{p}}^B$, in which case $\boldsymbol{\omega}^{\hat{b}}_{t+1} = (\mathbb{N}^{A^{\hat{b}}}\backslash i, (\mathbf{x}^b, \hat{p}^{\hat{b}}), \hat{p}^{\hat{b}})$ and for $b \neq \hat{b}$, $\boldsymbol{\omega}^b_{t+1} = (\mathbb{N}^{A^b}, \mathbf{x}^b, \hat{p}^b)$, or the buyer clock prices reach $\tilde{\mathbf{p}}^B$ such that $\min_{b\in\mathbb{Z}^B} \phi^b(\tilde{p}^b,\boldsymbol{\omega}_t) = \tau(\boldsymbol{\omega}_t)$ with no exit, in which case for all b, $\boldsymbol{\omega}^b_{t+1} = (\mathbb{N}^{A^b}, \mathbf{x}^b, \tilde{p}^b)$.

Decrease the vector of seller clock prices from $\boldsymbol{p}^S = (p^s)_{s\in\mathbb{Z}^S}$ by decreasing the clocks for the smallest number of seller groups possible so as to decrease $\max_{s\in\mathbb{Z}^S} \gamma^s(\cdot,\boldsymbol{\omega}_t)$ (or by decreasing the clock of only one seller group if $\max_{s\in\mathbb{Z}^S} \gamma^s(\cdot,\boldsymbol{\omega}_t)$ is nonincreasing) until either there is an exit by a group $\hat{s}$ seller j at clock price vector $\hat{\mathbf{p}}^S$, in which case $\boldsymbol{\omega}^{\hat{s}}_{t+1} = (\mathbb{M}^{A^{\hat{s}}}\backslash j, (\mathbf{x}^s, \hat{p}^{\hat{s}}), \hat{p}^{\hat{s}})$ and for $s \neq \hat{s}$, $\boldsymbol{\omega}^s_{t+1} = (\mathbb{M}^{A^s}, \mathbf{x}^s, \hat{p}^s)$, or the seller clock prices reach $\tilde{\mathbf{p}}^S$ such that $\max_{s\in\mathbb{Z}^S} \gamma^s(\tilde{p}^s,\boldsymbol{\omega}_t) = \tau(\boldsymbol{\omega}_t)$ with no exit, in which case for all s, $\boldsymbol{\omega}^s_{t+1} = (\mathbb{M}^{A^s}, \mathbf{x}^s, \tilde{p}^s)$.

If both targets are reached with no exit, then $e_{t+1} = 1$; otherwise $e_{t+1} = 0$.

If $\sum_{b\in\mathbb{Z}^B} n^{A^b} > \sum_{s\in\mathbb{Z}^S} m^{A^s}$, increase only the buyer clock prices as above until there is an exit, and similarly for sellers if $\sum_{b\in\mathbb{Z}^B} n^{A^b} < \sum_{s\in\mathbb{Z}^S} m^{A^s}$. The states transition analogously to the symmetric case.

Group BOPP mechanism

Given the numbers of active buyers and sellers in each group, $\mathbf{j} \equiv (j^1, ..., j^{z^B}, \hat{\jmath}^1, ..., \hat{\jmath}^{z^S})$ (recall that z^B is the number of buyer groups and z^S is the number of seller groups) and letting $I_{\mathbf{j}} \equiv [\min_{b \text{ s.t. } j^b>0} \Phi^b_\alpha(p^b), \max_{s \text{ s.t. } \hat{\jmath}^s>0} \Gamma^s_\alpha(p^s)]$, define

$$\hat{\theta}^*_{\mathbf{j}} \equiv \begin{cases} \max_{s \text{ s.t. } \hat{\jmath}^s>0} \Gamma^s_\alpha(p^s), & \text{if } \hat{J}(\theta) < 0 \text{ for all } \theta \in I_{\mathbf{j}} \\ \hat{\theta}, & \text{if } \hat{J}(\hat{\theta}) = 0 \text{ for some } \hat{\theta} \in I_{\mathbf{j}} \\ \min_{b \text{ s.t. } j^b>0} \Phi^b_\alpha(p^b), & \text{if } \hat{J}(\theta) > 0 \text{ for all } \theta \in I_{\mathbf{j}}, \end{cases} \tag{5}$$

where

$$\hat{J}(\theta) \equiv \sum_{b\in\mathbb{Z}^B} j^b \frac{1 - F^b(\Phi^{b^{-1}}_\alpha(\theta))}{f^b\left(\Phi^{b^{-1}}_\alpha(\theta)\right)\Phi^{b^{-1}\prime}_\alpha(\theta)} - \sum_{s\in\mathbb{Z}^S} \hat{\jmath}^s \frac{G^s(\Gamma^{s^{-1}}_\alpha(\theta))}{g^s\left(\Gamma^{s^{-1}}_\alpha(\theta)\right)\Gamma^{s^{-1}\prime}_\alpha(\theta)}.$$

Proposition 9 characterizes the **group Bayesian optimal privacy preserving (G-BOPP)** mechanism, which based on types of agents who do not trade, maximizes the probability that $n^A = m^A$ trades occur, conditional on n^A trades being optimal.

Proposition 9 *Any G-BOPP mechanism is characterized by the target function* $\tau(\boldsymbol{\omega}_t) = \theta^*_{(\mathbf{n}^A,\mathbf{m}^A)}$*, where* $\mathbf{n}^A \equiv (n^{A^1}, ..., n^{A^{z^B}})$, $\mathbf{m}^A \equiv (m^{A^1}, ..., m^{A^{z^S}})$*, and* $\theta^*_{(\mathbf{n}^A,\mathbf{m}^A)}$ *is defined in*

(5), *and comparison functions*

$$\phi^b(p^b,\boldsymbol{\omega}_t)=\begin{cases}\Phi^b_\alpha(p^b), & \text{if } n^{A^b}>0\\ \infty, & \text{otherwise}\end{cases} \quad \text{and } \gamma^s(p^s,\boldsymbol{\omega}_t)=\begin{cases}\Gamma^s_\alpha(p^s), & \text{if } m^{A^s}>0\\ -\infty, & \text{otherwise.}\end{cases}$$

Group augmented prior-free mechanism

We now construct the **group augmented prior-free mechanism**, which is a prior-free approximation to the G-BOPP mechanism using comparison and target functions given by

$$\phi^b(p^b,\boldsymbol{\omega}_t)=\begin{cases}\tilde{\Phi}^b_\alpha(n^{A^b}+1;p^b), & \text{if } n^{A^b}>0\\ \infty, & \text{otherwise,}\end{cases} \quad \text{and } \gamma^b(p^s,\boldsymbol{\omega}_t)=\begin{cases}\tilde{\Gamma}^s_\alpha(m^{A^s}+1;p^S), & \text{if } m^{A^s}>0\\ -\infty, & \text{otherwise,}\end{cases}$$

and $\tau(\boldsymbol{\omega}_t)=\tilde{\theta}^*$ defined analogously to above, with an interior solution satisfying

$$\max_{b\in\mathbb{Z}^B \text{ s.t. } \phi^b(p^b,\boldsymbol{\omega}_t)\leq\tilde{\theta}^*} \frac{\tilde{\theta}^*-\tilde{\Phi}^b_\alpha(n^{A^b}+1;p^b)}{\sigma^b_{j+1}} = \max_{s\in\mathbb{Z}^S \text{ s.t. } \gamma^b(p^s,\boldsymbol{\omega}_t)\geq\tilde{\theta}^*} \frac{\tilde{\Gamma}^s_\alpha(m^{A^s}+1;p^s)-\tilde{\theta}^*}{\sigma^s_{j+1}}.$$

In words, the target virtual type $\tilde{\theta}^*$ is chosen to equalize the maximum changes required in the smoothed weighted virtual values and in the smoothed weighted virtual costs (adjusted for the estimated change in virtual types between order statistics) in order to have those virtual values above the virtual costs for all the active buyers and sellers.

5.3 Results

Using the arguments in the proof of Proposition 2, one can show that the results on asymptotic optimality extend to heterogeneous groups if the number of buyers and sellers in each group goes to infinity in fixed proportions: for all $b\in\mathbb{Z}^B$ and $s\in\mathbb{Z}^S$, $\lim_{n,m\to\infty}\frac{n^b}{n}$ and $\lim_{n,m\to\infty}\frac{m^s}{n}$ are positive and finite.

With the qualification that with heterogeneous groups a mechanism is envy free if there is no envy within groups,[28] the following corollary summarizes the properties of the group augmented prior-free mechanism:

Corollary 1 *The group augmented prior-free mechanism is asymptotically optimal, ex post individually rational, dominant strategy incentive compatible, non-bossy, envy free,*

[28] Of course, because the allocation is based on virtual types, whose ranking will not coincide with the ranking of true types with asymmetric groups, buyers or sellers from one group may envy those from another.

and deficit free. It permits an implementation via a multi-clock auction that is self-enforcing.

6 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 comparison and target functions for a double or multi-clock auction. Such constraints may restrict the ability of members of particular subsets of agents to trade, or consist of minimum revenue thresholds. Alternative objectives could include the maximization of the weighted objective subject to achieving an ex post revenue threshold or favoritism towards subsets of agents. We now briefly discuss each of these modifications.

6.1 Real-time diagnostics and ex post revenue thresholds

In practice, mechanism designers may face "hard" revenue constraints ex post. For example, the FCC's incentive auction faces a revenue constraint.[29] In a prior-free setup, there is of course no guarantee that generating a given amount of revenue is feasible in expectation, let alone ex post. To avoid the disastrous outcome in which the double-clock auction continues until no trade occurs, the mechanism needs to provide real-time feedback as to whether achieving the revenue threshold is feasible. If it is not feasible, then it is certainly preferable to stop the auction when the maximum revenue is achieved. An advantage our double-clock and multi-clock auctions is that they permit the provision of such real-time diagnostics.

As an illustration, consider the symmetric case with values and costs drawn from the Uniform distribution as displayed in Table 2. If the designer's revenue threshold is, for example, 0.5, the double-clock auction would end with four active buyers and sellers and prices $p^B = 0.59$ and $p^S = 0.25$. In this case, the revenue threshold is feasible ex post. In contrast, if the threshold were, say, 1.2, the double-clock auction would end with three trades and a revenue of 1.02 because with three active pairs $\tilde{\Phi}_1$ exceeds $\tilde{\Gamma}_1$, indicating that revenue decreases if the auction continues. Although this outcome fails to achieve the revenue threshold ex post, it is preferred to continuing the auction both with regards

[29] As stated by Milgrom and Segal (2015, p.42), "In the incentive auction, revenues for the forward auction portion must be sufficient to pay the costs the broadcasters incur in moving to new broadcast channels, as well as meeting certain other gross and net revenue goals, so the possibility of including a cost target in the scoring rule is necessary to make the whole design feasible."

n^A	p^B	p^S	Revenue	$\tilde{\Phi}_1(n^A+1)$ $\geq \tilde{\Gamma}_1(n^A+1)$	Threshold α	Est. Δ in social surplus if continue	Est. Δ in revenue if continue	Actual Δ in revenue if continue
10	0.	1.	−10.					
9	0.16	0.98	−7.33	No		0.63	2.27	2.27
8	0.24	0.87	−5.05	No		0.45	1.91	0.93
7	0.28	0.87	−4.13	No		0.48	1.26	3.67
6	0.43	0.5	−0.46	No	0.12	−0.17	1.31	0.48
5	0.45	0.44	0.03	No	0.2	−0.22	0.84	0.53
4	0.49	0.35	0.56	No	0.39	−0.38	0.59	0.47
3	0.59	0.25	1.02	Yes	1.15	−0.48	−0.06	−0.16
2	0.6	0.17	0.86	Yes	2.02	−0.57	−0.29	−0.17
1	0.86	0.17	0.7	Yes	4.75	−0.88	−0.7	−0.7
0	0.92	0.03	0.	Yes				

Table 2: Real-time diagnostics for a double-clock auction with 10 buyer-seller pairs with values and costs drawn from $U[0,1]$.

to the revenue and social surplus that is generated.

With slight modifications, the mechanism can also be extended to account for a hybrid objective function that consists of first achieving a revenue target and, conditional on this being achieved, puts weight α on revenue and weight $1-\alpha$ on social surplus.

6.2 Alternative objectives and additional constraints

Caps and group-specific revenue constraints

A multi-clock auction can incorporate other constraints too. For example, a designer or regulator may want to cap the number of units a subset of buyers acquires. Within a multi-clock auction, this can be done by treating that set of bidders as a group and starting the procedure by advancing the clock for that group until the number of active agents in the group is reduced to the number eligible to trade. This is implemented by letting that group's comparison function be $-\infty$ if they are buyers, or ∞ if they are sellers, as long as the number of active members of the group remains above the cutoff.

The multi-clock auction can also accommodate the requirement that members of some buyer group $\hat{b}$ contribute payments of at least $\underline{R}$ in order for any members of that group to trade, in which case the designer adjusts the comparison function for group $\hat{b}$ buyers to be

$$\phi^{\hat{b}}(p^{\hat{b}}, \boldsymbol{\omega}_t) = \begin{cases} -\infty, & \text{if } n^{A^{\hat{b}}} > 0 \text{ and } \underline{R} < p^{\hat{b}} n^{A^{\hat{b}}} \\ \tilde{\Phi}^{\hat{b}}_{\alpha}(n^{A^{\hat{b}}}+1; p^{\hat{b}}), & \text{if } n^{A^{\hat{b}}} > 0 \text{ and } \underline{R} \geq p^{\hat{b}} n^{A^{\hat{b}}} \\ \infty, & \text{otherwise.} \end{cases}$$

Favoring groups of agents

The multi-clock auction is also flexible enough to allow the designer to favor a particular subset of agents over other agents.[30] This can be accomplished using a multi-clock auction that assigns favored and nonfavored agents to separate groups, evaluates nonfavored agents using virtual types that incorporate the designer's unconstrained weight α on revenue, and evaluates favored agents using virtual types with weight α^f on revenue. Favoritism then simply means $\alpha^f < \alpha$. In the multi-clock auction, nonfavored buyers have comparison function defined in terms of $\tilde{\Phi}_\alpha$, and favored buyers have comparison function defined in terms of $\tilde{\Phi}_{\alpha^f}$, and similarly for sellers. The target function is defined as before, but uses the different weights on revenue for the different groups.

6.3 Equivalence to full-information first-price auctions

The algorithmic structure and pricing rule of the augmented prior-free mechanism means that the outcome of the dominant strategy equilibrium in this mechanism has a payoff equivalent outcome in a first-price double auction in which agents are completely informed about each others' types. Thus, under complete information, we have revenue equivalence between a first-price double auction and our second-price double auction. Proposition 6 in Milgrom and Segal (2015), which inspired this thought experiment, provides a similar equivalence result for a one-sided auction.

To elaborate only briefly, consider the symmetric setup in which virtual types are increasing. Let the type realizations be such that the BOPP mechanism induces the quantity traded k, in which case prices are $v_{(k+1)}$ and $c_{[k+1]}$ if $\Phi_\alpha(v_{(k+1)}) \geq \Gamma_\alpha(c_{[k+1]})$ and otherwise prices are $\Phi_\alpha^{-1}(\theta_k^*)$ and $\Gamma_\alpha^{-1}(\theta_k^*)$, where $\theta_k^* \in [\Phi_\alpha(v_{(k+1)}), \Gamma_\alpha(c_{[k+1]})]$ is as defined in Proposition 5. The full-information first-price auction with the corresponding allocation rule has an equilibrium in which all losing traders bid their types and all trading agents bid the BOPP prices.

In the prior-free setup when distributions are not known by any participants or the designer, the designer announces the allocation rule corresponding to the augmented prior-free mechanism, which is a function of the losing bids. Agents who lose still bid their types in this equilibrium, and agents who trade submit the minimal or maximal bids that allow them to trade, given the allocation rule.

[30] 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))

6.4 Nonparametric versus parametric approaches

Our nonparametric approach for estimation works well when the underlying Bayesian design problem does not require ironing, but not necessarily otherwise because in a nonparametric approach one cannot detect nonmonotonicities that occur after the first time the weighted virtual type functions cross. In contrast, if parametric assumptions are imposed on the distributions F and G, then one could make out-of-sample predictions for the virtual types of the inframarginal agents. This would allow the mechanism to continue decreasing the quantity traded even after the first point of intersection of the estimated virtual types if the estimated parameters indicates that the underlying problem is nonmonotone and requires ironing.

In this sense, the question of whether a unique local maximum is a reasonable assumption is intertwined with the question of whether one is willing to impose some parametric restrictions on the distributions. For example, one could follow the approach in Loertscher and Niedermayer (2015), according to which the virtual type functions are modelled as Chebychev polynomials. As noted by Loertscher and Niedermayer, this approach has the benefits that it is flexible, permits testing for regularity with a simple parameter restriction, and by incorporating sufficiently many polynomial terms, permits accurate extrapolation. Whether the assumption of a unique local maximum or of a particular parameterization is more restrictive is, obviously, an empirical question. At least in principle, it could be answered based on practical implementation of the augmented prior-free mechanism as a direct mechanism because that would allow one to detect the need for ironing ex post, that is, after the mechanism has been run and trades and payments have occurred.

7 Conclusions

We develop a prior-free double-clock auction that is asymptotically optimal, endows agents with obviously dominant strategies to bid truthfully, preserves the privacy of trading agents, and performs well relative to the optimal mechanism even when the number of buyers and sellers is small. The mechanism has the additional desirable properties that it never runs a deficit, is envy free and weakly group strategy-proof, requires only limited commitment by the designer, respects agents' individual rationality constraints ex post, and its equilibrium outcome remains an equilibrium outcome in a full-information first-price double auction. Properties such as these seem important for market design to move forward towards practical implementation in environments in which both buyers and sellers are privately informed about their types.

Methodologically, we exploit the connection between the empirical measure of spacings of order statistics and the theoretical construct of virtual types. Although the direct empirical counterpart to virtual types based on one spacing is too volatile to provide useful guidance, we show that excessive volatility can be overcome through appropriate smoothing of spacings without any cost for asymptotic optimality.

Privacy preservation restricts what can be achieved even in a Bayesian setup where distributions are known by the designer. We establish the impossibility of ex post efficient privacy preserving trade under the mild condition that full trade is sometimes but not always optimal and provide a Bayesian benchmark for what can optimally be achieved subject to privacy preservation when the designer knows the distributions.

Many features of the mechanisms we develop, such as prior-freeness, real-time diagnostics, and the flexibility to accommodate various constraints and to pursue a combination of revenue and surplus goals, may prove useful in a number of setups and applications, including internet platforms and one-sided allocation problems.

A Appendix: Proofs

Proof of Lemma 1. We show the result for costs. The result for values follows analogously. In this proof and those that follow, where we use notation such as ρm as an index, e.g., $c_{[\rho m]}$, we mean the greatest integer less than or equal to ρm. Recall that $\chi_\alpha(j) \equiv \max\{0,\ \alpha(j-2)-(1-\alpha)\}$. Because $\chi_0(\rho m)=0$, the result holds for $\alpha=0$, so assume $\alpha>0$. Because $\overline{\rho}<1$, we are away from the boundary and can focus on $\sigma^c_{\rho m}=\frac{c_{[\rho m+r_m]}-c_{[\rho m]}}{r_m}$ ($\sigma^c_{\rho m}$ is defined differently close to the boundary, i.e., for $\rho m > m-r_m$). To show that $V(\chi_\alpha(\rho m)\sigma^c_{\rho m})$ goes to zero with m, it is sufficient (given assumptions of continuity on a bounded support) to show that $\chi_\alpha(\rho m)^2 V\left[\frac{G(c_{[\rho m+r_m]})-G(c_{[\rho m]})}{r_m}\right]$ goes to zero, and because the cdf of an order statistic is itself a uniform order statistic, it is sufficient to show that $\chi_\alpha(\rho m)^2 V\left[\frac{u_{[\rho m+r_m]}-u_{[\rho m]}}{r_m}\right]$ goes to zero, where $u_{[i]}$ is the ith order statistic out of m draws from $U[0,1]$.

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

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

It follows that $\chi_\alpha(\rho m)^2 V\left[\frac{u_{[\rho m+r_m]}-u_{[\rho m]}}{r_m}\right]=\chi_\alpha(\rho m)^2\frac{m+1-r_m}{(m+1)^2(m+2)r_m}$, which is nondecreasing in ρ. Taking the limit of the above expression, we have $\lim_{m\to\infty}\chi_\alpha(\rho m)^2\frac{m+1-r_m}{(m+1)^2(m+2)r_m}=$

$\lim_{m\to\infty} \frac{\chi_\alpha(\rho m)^2}{m^2}\frac{1}{r_m} = 0$, where the first equality uses $\lim_{m\to\infty} \frac{r_m}{m} = 0$ and the second equality uses $\lim_{m\to\infty} r_m = \infty$ and the fact that $\chi_\alpha(\rho m)$ is of order m. This gives us the existence of a uniform bound. Q.E.D.

Proof of Lemma 2. Again, we show the result for virtual costs, with the result for virtual values following analogously. Also as above, where we use ρm as an index, we mean the greatest integer less than or equal to ρm. Three auxiliary results will be useful. First, for $j \in \{1, ..., m-1\}$, using the fact that the density of the jth 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 follows that

$$\begin{aligned}
E_{\mathbf{c}}\left[\frac{G(c_{[j]})}{g(c_{[j]})}\right] &= \int_{\underline{c}}^{\overline{c}} \frac{G(x)}{g(x)}\frac{m!}{(j-1)!(m-j)!}G^{j-1}(x)(1-G(x))^{m-j}g(x)dx \\
&= \int_{\underline{c}}^{\overline{c}} \frac{m!}{(j-1)!(m-j)!}G^{j}(x)(1-G(x))^{m-j}dx \\
&= (m-j)\int_{\underline{c}}^{\overline{c}} \frac{m!}{(j-1)!(m-j)!}xG^{j}(x)(1-G(x))^{m-j-1}g(x)dx \\
&\quad -j\int_{\underline{c}}^{\overline{c}} \frac{m!}{(j-1)!(m-j)!}xG^{j-1}(x)(1-G(x))^{m-j}g(x)dx \qquad (6)\\
&= j\int_{\underline{c}}^{\overline{c}} \frac{m!}{j!(m-j-1)!}xG^{j}(x)(1-G(x))^{m-j-1}g(x)dx \\
&\quad -j\int_{\underline{c}}^{\overline{c}} \frac{m!}{(j-1)!(m-j)!}xG^{j-1}(x)(1-G(x))^{m-j}g(x)dx \\
&= jE_{\mathbf{c}}[c_{[j+1]} - c_{[j]}],
\end{aligned}$$

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.

Second, one can show that for $\rho \in (0,1)$,

$$\lim_{m\to\infty} \rho m E_{\mathbf{c}}\left[c_{[\rho m+1]} - c_{[\rho m]} - \frac{c_{[\rho m + r_m]} - c_{[\rho m]}}{r_m}\right] = 0. \qquad (7)$$

To see this, note that given $\rho \in (0,1)$ and m sufficiently large, $\rho m + r_m \leq m$, so the expression in (7) is well defined. We can then write the expression inside the limit in (7) as

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

where the first equality writes $c_{[\rho m+r_m]}-c_{[\rho m]}$ as the sum of r_m spacings, the second equality uses (6), the third equality uses the fact that a cdf evaluated at an order statistic is itself a uniform order statistic, with $u_{[j]}$ denoting the jth order statistic out of m draws from $U[0,1]$, the fourth equality rearranges, and the fifth equality uses $E_{\mathbf{u}}\left[u_{[j]}\right]=\frac{j}{m+1}$. Taking the limit of the above expression as m goes to infinity, we get

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

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

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

$$
\lim_{m\to\infty}\sup_{\rho\in[0,\overline{\rho}]}V_{\mathbf{c}}\left[\frac{G(c_{[\rho m]})}{g(c_{[\rho m]})}\right]=0. \tag{8}
$$

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

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

Taking the limit, we have

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

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

With these three auxiliary results in hand, we now proceed with the proof of Lemma 2. Denote the difference between theoretical and smoothed virtual costs as

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

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

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

where the final equality uses (6).

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

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

Taking the limit, we have

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

where the first equality uses $\lim_{m\to\infty} E_{\mathbf{c}}\left[\frac{c_{[\rho m+r_m]}-c_{[\rho m]}}{r_m}\right] = 0$ (the numerator is bounded by

$\overline{c} - \underline{c}$ and r_m goes to infinity with m) and the second equality uses (7).

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

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

By (8), $\lim_{m\to\infty}\sup_{\rho\in[0,\overline{\rho}]} V_{\mathbf{c}}\left[\frac{G(c_{[\rho m]})}{g(c_{[\rho m]})}\right] = 0$. By Lemma 1, for m sufficiently large, $V_{\mathbf{c}}[\chi_\alpha(\rho m)\sigma^c_{\rho m}]$ $\leq u^c(\rho, m) \leq u^c(\overline{\rho}, m)$, where $\lim_{m\to\infty} u^c(\overline{\rho}, m) = 0$. Thus, $\lim_{m\to\infty}\sup_{\rho\in[0,\overline{\rho}]} V_{\mathbf{c}}\left[\chi_\alpha(\rho m)\sigma^c_{\rho m}\right] =$ 0. It follows then by the Cauchy-Schwarz inequality, that the limit of the covariance terms is also zero. Thus,

$$\lim_{m\to\infty}\sup_{\rho\in[0,\overline{\rho}]} V_{\mathbf{c}}\left[Y_{\rho m}\right] = 0. \tag{9}$$

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

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

where the first equality uses $\lim_{m\to\infty} E_{\mathbf{c}}\left[Y_{\rho m}\right] = 0$ and the second equality uses (9). This establishes uniform convergence in probability to zero. Q.E.D.

Proof of Lemma 3. By the assumption of a unique local maximum, in the limit as n and m go to infinity, the optimal mechanism is a price posting mechanism such that buyers are charged the price v^* and trading sellers are paid c^*, with cutoff types v^* and c^* defined by $\Phi_\alpha(v^*) = \Gamma_\alpha(c^*)$ and $n(1 - F(v^*)) = mG(c^*)$ (see, for example, Loertscher and Niedermayer (2013)). If $\alpha = 0$, then the result follows from McAfee (1992), so we consider the case with $\alpha > 0$, which together with our assumption that $\underline{v} \leq \overline{c}$, implies that $v^* \in (\underline{v}, \overline{v})$ and $c^* \in (\underline{c}, \overline{c})$. It follows that there exists $\overline{\rho} \in (0, 1)$ sufficiently large such that $v^* \in (\overline{v} - \overline{\rho}(\overline{v} - \underline{v}), \overline{v})$ and $c^* \in (\underline{c}, \underline{c} + \overline{\rho}(\overline{c} - \underline{c}))$, which will allow us to use Lemma 2.

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

$$\lim_{\varepsilon\to 0} v^\varepsilon = v^*, \ \lim_{\varepsilon\to 0} v^{-\varepsilon} = v^*, \ \lim_{\varepsilon\to 0} c^\varepsilon = c^*, \text{ and } \lim_{\varepsilon\to 0} c^{-\varepsilon} = c^*. \tag{10}$$

This is illustrated in Figure 4 (note that values are decreasing and costs are increasing along the horizontal axis).

By Lemma 2, for all $\varepsilon > 0$, $\lim_{m\to\infty}\Pr\left(\sup_{\rho\in[0,\overline{\rho}]}\left|\Gamma_\alpha(c_{[\rho m]}) - \tilde{\Gamma}_\alpha(\rho m)\right| \geq \varepsilon\right) = 0$, which

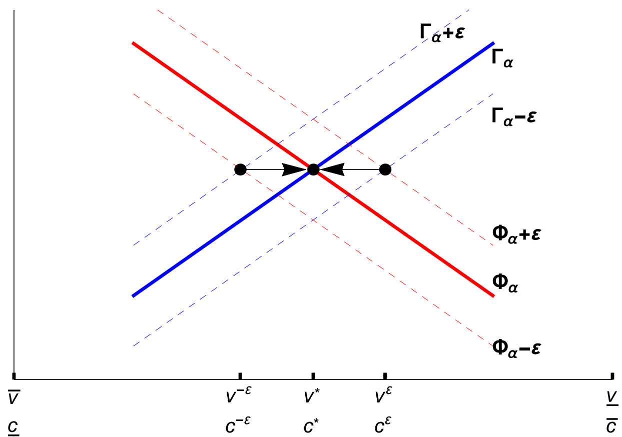

Figure 4: Illustration of $v^\varepsilon < v^* < v^{-\varepsilon}$ and $c^{-\varepsilon} < c^* < c^\varepsilon$ and convergence to v^* and c^*.

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

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

Thus, using (10), for all $\varepsilon > 0$,

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

which completes the proof. Q.E.D.

Proof of Proposition 3. That double-clock auctions satisfy dominant strategies (DS), non-bossiness (NB), and envy-freeness (EF) is obvious. Therefore, we show that DS, NB, and EF permit implementation via double-clock auction. DS implies that every agent faces a price that does not depend on his report, and EF implies that all buyers and all sellers face the same prices. NB implies that these prices do not vary with the reports of agents who trade, and hence can only be functions of reports of agents who do not trade. DS and EF imply that the buyers' price must increase and the sellers' price must decrease as the quantity traded decreases. This permits a double-clock auction implementation: Assume that the quantity traded in the direct mechanism satisfying DS, NB, and EF when the reports are $(\boldsymbol{v}, \boldsymbol{c})$ is q. Because of DS and NB, the buyers' price p^B can be written as $p^B = p^B(v_{(q+1)}, .., v_{(n)}, c_{[q+1]}, .., c_{(m)})$ and the sellers' price p^S as $p^S = p^S(v_{(q+1)}, .., v_{(n)}, c_{[q+1]}, .., c_{[m]})$. In the double-clock auction, $v_{(n)}, .., v_{(q+1)}$ and

$c_{[m]},..,c_{[q+1]}$ would be the exit prices of the $n-q$ buyers with smallest values and the $m-q$ sellers with the largest costs. The designer can end the procedure when the clock prices reach $p^B(v_{(q+1)},..,v_{(n)},c_{[q+1]},..,c_{(m)})$ and $p^S(v_{(q+1)},..,v_{(n)},c_{[q+1]},..,c_{(m)})$, establishing that a double-clock auction implementation exists.

We omit the proof that a double-clock auction satisfies privacy preservation and endows agents with obviously dominant strategies. We show that it satisfies weak group strategy-proofness. A bidder that exits when the clock price reaches its type achieves a payoff of at least zero. Consider the first clock prices at which the double-clock auction is affected by a group deviation. A bidder cannot benefit by a deviation that results in its losing and getting a deviation payoff of zero, so the first deviation must be made by a buyer who chooses not to exit when the buyer clock price reaches his value, and who becomes winning, or by a seller who chooses not to exit when the seller clock price reaches his cost, and who becomes winning. In the case of a buyer, the buyer's price must then be greater than his value, so that buyer loses from the group deviation, and in the case of a seller, the seller's price must then be less than his cost, so that seller loses from the group deviation. Q.E.D.

Proof of Proposition 5. For a mechanism to be the BOPP mechanism, we require that θ be chosen to maximize the probability of achieving the n^Ath trade. This probability is $\left(1-H_{n^A}\left(\Phi_\alpha^{-1}(\theta)\right)\right)L_j\left(\Gamma_\alpha^{-1}(\theta)\right)$, where $H_{n^A}(v)$ is the probability that the n^Ath highest draw of a buyer is no more than v conditional on the n^A+1st highest draw and $L_{n^A}(c)$ is the probability that the n^Ath lowest cost draw is no more than c conditional on the n^A+1st lowest draw. Denote the associated densities by h_{n^A} and ℓ_{n^A}. The first-order condition is

$$\frac{1-H_{n^A}(\Phi_\alpha^{-1}(\theta))}{h_{n^A}(\Phi_\alpha^{-1}(\theta))}\frac{1}{\Phi_\alpha^{-1\prime}(\theta)}=\frac{L_{n^A}(\Gamma_\alpha^{-1}(\theta))}{\ell_{n^A}(\Gamma_\alpha^{-1}(\theta))}\frac{1}{\Gamma_\alpha^{-1\prime}(\theta)}. \tag{11}$$

Using $1-H_{n^A}(v)=\left(\frac{1-F(v)}{1-F(v_{(n^A+1)})}\right)^{n^A}$ for $v>v_{(n^A+1)}$ and $L_{n^A}(c)=\left(\frac{G(c)}{G(c_{[n^A+1]})}\right)^{n^A}$ for $c<c_{[n^A+1]}$, the result follows. Q.E.D.

Proof of Proposition 6. The first two statements of the proposition follow from Propositions 1 and 3. Suppose $\tilde{\Phi}_\alpha(j)<\tilde{\Gamma}_\alpha(j)$. If $\hat{v}_{(j-1)}<\hat{c}_{[j-1]}$, then trade at prices $\hat{v}_{(j-1)}$ and $\hat{c}_{[j-1]}$ violates the constraint not to run a deficit, so suppose $\hat{v}_{(j-1)}\geq\hat{c}_{[j-1]}$. The incremental weighted objective from the trade of $j-1$ units at p^B and p^S rather than

$j-2$ units at the predicted prices $\hat{v}_{(j-1)} = v_{(j)} + \sigma_j^v$ and $\hat{c}_{[j-1]} = c_{[j]} - \sigma_j^c$ is

$$\begin{aligned} & \alpha(j-1)(p^B - p^S) - \alpha(j-2)\left(v_{(j)} + \sigma_j^v - c_{[j]} + \sigma_j^c\right) + (1-\alpha)\left(v_{(j)} + \sigma_j^v - c_{[j]} + \sigma_j^c\right) \\ = & \begin{cases} \alpha(j-1)\left(\tilde{\Phi}_\alpha(j;p^B) - \tilde{\Gamma}_\alpha(j;p^S)\right) - \chi_\alpha(j)\left(\tilde{\Phi}_\alpha(j) - \tilde{\Gamma}_\alpha(j)\right), & \text{if } j \geq \frac{1+\alpha}{\alpha} \\ \alpha(j-1)(p^B - p^S) - (\alpha(j-1)-1)\left(\hat{v}_{(j-1)} - \hat{c}_{[j-1]}\right), & \text{otherwise,} \end{cases} \end{aligned}$$

which is nonnegative if $\tilde{\Phi}_\alpha(j;p^B) \geq \tilde{\Gamma}_\alpha(j;p^S)$ given the suppositions that $\tilde{\Phi}_\alpha(j) < \tilde{\Gamma}_\alpha(j)$ and $\hat{v}_{(j-1)} \geq \hat{c}_{[j-1]}$ (note that $j < \frac{1+\alpha}{\alpha}$ implies that $\alpha(j-1)-1 < 0$). Q.E.D.

Proof of Proposition 7. Fix $\min\{n,m\}$. As $\overline{v}$ increases, the probability that $\min\{n,m\}$ trades is efficient goes to one. In the fixed-target BOPP mechanism, $\theta_{\min\{n,m\}} = 1/2$, which implies that $\min\{n,m\}$ trades occur if and only if the $\min\{n,m\}$ lowest costs are less than $1/2$ and the $\min\{n,m\}$ highest values are greater than $1/2$, which occurs with probability $G^{\min\{n,m\}}(\frac{1}{2})(1-F(\frac{1}{2}/\overline{v}))^{\min\{n,m\}}$, which (using $G(\frac{1}{2}) = \frac{1}{2}$ by symmetry) goes to $\frac{1}{2^{\min\{n,m\}}}$ as $\overline{v}$ increases. In the BOPP mechanism, as $\overline{v}$ increases $\theta^*_{\min\{n,m\}}$ goes to 1, which implies that $\min\{n,m\}$ trades occur with probability $(1-F(\frac{1}{2}/\overline{v}))^{\min\{n,m\}}$, which goes to 1 as $\overline{v}$ increases. Q.E.D.

Proof of Proposition 8. For a given type realization, if the BOPP mechanism has $k > 0$ trades, then $v_{(k)} \geq c_{[k]}$, which implies that $\tilde{\Phi}_0(k) \geq \tilde{\Gamma}_0(k)$, and so the augmented prior-free mechanism has at least $k-1$ trades. For the second part of the proposition, assume $\underline{v} < \overline{c}$. The augmented prior-free mechanism has $\min\{n,m\}-1$ trades when $v_{(\min\{n,m\})} < c_{[\min\{n,m\}]}$ and $c_{[\min\{n,m\}-1]} \leq (v_{(\min\{n,m\})} + c_{[\min\{n,m\}]})/2 \leq v_{(\min\{n,m\}-1)}$, which occurs with positive probability. When, in addition, it is not the case that $c_{[\min\{n,m\}-1]} \leq \theta^*_{\min\{n,m\}-1} \leq v_{(\min\{n,m\}-1)}$, the BOPP mechanism has only $\min\{n,m\}-2$ trades. Because $\theta^*_{\min\{n,m\}-2}$ does not depend on the type realizations, except for bounds, it differs from $(v_{(\min\{n,m\})} + c_{[\min\{n,m\}]})/2$ with probability one, and so with positive probability the stated conditions hold. Q.E.D.

Proof of Proposition 9. The optimization problem is to choose θ to maximize

$$\prod_{b \in \hat{\mathbb{Z}}^B} \left(1 - H^b_{n^{A^b}}(\Phi_\alpha^{b^{-1}}(\theta))\right) \prod_{s \in \hat{\mathbb{Z}}^S} L^s_{m^{A^s}}\left(\Gamma_\alpha^{s^{-1}}(\theta)\right).$$

The first-order condition can be written as

$$\textstyle\sum_{b \in \hat{\mathbb{Z}}^B} \frac{h^b_{n^{A^b}}(\Phi_\alpha^{b^{-1}}(\theta))}{1 - H^b_{n^{A^b}}(\Phi_\alpha^{b^{-1}}(\theta))} \Phi_\alpha^{b^{-1}\prime}(\theta) = \sum_{s \in \hat{\mathbb{Z}}^S} \frac{\ell^s_{m^{A^s}}\left(\Gamma_\alpha^{s^{-1}}(\theta)\right)}{L^s_{m^{A^s}}\left(\Gamma_\alpha^{s^{-1}}(\theta)\right)} \Gamma_\alpha^{s^{-1}\prime}(\theta).$$

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

References

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

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

Ausubel, L. (2004): "An Efficient Ascending-Bid Auction for Multiple Objects," *American Economic Review*, 94(5), 1452–1475.

Ausubel, L., and P. Milgrom (2006): "The Lovely But Lonely Vickrey Auction," in *Combinatorial Auctions*, ed. by P. Cramton, Y. Shoham, and R. Steinberg, pp. 17–40. Cambridge University Press, Cambridge and London.

Baliga, S., and R. Vohra (2003): "Market Research and Market Design," *Advances in Theoretical Economics*, 3(1), Article 5.

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

Bergemann, D., and S. Morris (2005): "Robust Mechanism Design," *Econometrica*, 73(6), 1771–1813.

——— (2009): "Robust Implementation in Direct Mechanisms," *Review of Economic Studies*, 76(4), 1175–1204.

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

Brandt, F., and T. Sandholm (2008): "On the Existence of Unconditionally Privacy-Preserving Auction Protocols," *ACM Transactions on Information and Systems Security*, 11(2), 10:1–10:21.

Bulow, J., and J. Roberts (1989): "The Simple Economics of Optimal Auctions," *Journal of Political Economics*, 97(5), 1060–90.

Chen, N., N. Gravin, and P. Lu (2014): "Optimal Competitive Auctions," in *Proceedings of the 2014 ACM Symposium on Theory of Computing*, pp. 253–262. ACM.

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

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

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

Dütting, P., T. Roughgarden, and I. Talgam-Cohen (2014): "Modularity and Greed in Double Auctions," in *Proceedings of the Fifteenth ACM conference on Economics and Computation*, pp. 241–258. ACM.

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

Goldberg, A. V., J. D. Hartline, and A. Wright (2001): "Competitive Auctions and Digital Goods," in *Proceedings of the Twelfth Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 735–744. ACM/SIAM.

Gresik, T., and M. Satterthwaite (1989): "The Rate at which a Simple Market Converges to Efficiency as the Number of Traders Increases: An Asymptotic Result for Optimal Trading Mechanisms," *Journal of Economic Theory*, 48(1), 304–332.

Guerre, E., I. Perrigne, and Q. Vuong (2000): "Optimal Nonparametric Estimation of First-Price Auctions," *Econometrica*, 68(3), 525–574.

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

Hurwicz, L. (1973): "The Design of Mechanisms for Resource Allocation," *American Economic Review*, 63(2), 1–30.

Jackson, M. O., and I. Kremer (2007): "Envy-Freeness and Implementation in Large Economies," *Review of Economic Design*, 11(3), 185–198.

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

Li, S. (2015): "Obviously Strategy-Proof Mechanisms," Working Paper, Stanford University.

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

Loertscher, S., and C. Mezzetti (2015): "Two-Sided VCG and Double-Clock Auctions with Endogenous Reserve Prices," Working Paper, University of Melbourne.

Loertscher, S., and A. Niedermayer (2013): "Fee-Setting Mechanisms: On Optimal Intermediary Pricing and Transaction Taxes," Working Paper, University of Melbourne.

——— (2015): "Percentage Fees in Thin Markets: An Optimal Pricing Perspective," Working Paper, University of Mannheim and University of Melbourne.

Lucking-Reiley, D. (2000): "Vickrey Auctions in Practice: From Nineteenth-Century Philately to Twenty-First-Century E-Commerce," *Journal of Economic Perspectives*, 14(3), 183–192.

Marx, L. M., and J. M. Swinkels (1997): "Order Independence for Iterated Weak Dominance," *Games and Economic Behavior*, 18, 219–245.

McAfee, R. P. (1992): "A Dominant Strategy Double Auction," *Journal of Economic Theory*, 56(2), 434–450.

McMillan, J. (1994): "Selling Spectrum Rights," *Journal of Economic Perspectives*, 8(3), 145–162.

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

Myerson, R., and M. Satterthwaite (1983): "Efficient Mechanisms for Bilateral Trading," *Journal of Economic Theory*, 29(2), 265–281.

Myerson, R. B. (1981): "Optimal Auction Design," *Mathematical Operations Research*, 6(1), 58–73.

Norman, P. (2004): "Efficient Mechanisms for Public Goods with Use Exclusions," *Review of Economic Studies*, 71(4), 1163–1188.

Rohde, G. L. (2002): *Minority Broadcast Ownership*. Novinka.

Rothkopf, M. (2007): "Thirteen Reasons Why the Vickrey-Clarke-Groves Process Is Not Practical," *Operations Research*, 55, 191–197.

Rustichini, A., M. A. Satterthwaite, and S. R. Williams (1994): "Convergence to Efficiency in a Simple Market with Incomplete Information," *Econometrica*, 62(5), 1041–1063.

Satterthwaite, M., S. Williams, and K. Zachariadis (2015): "Price discovery Using a Double-Auction," SSRN Working Paper.

Satterthwaite, M., and S. R. Williams (1989): "The Rate of Convergence to Efficiency in the Buyer's Bid Double Auction as the Market Becomes Large," *Review of Economic Studies*, 56, 477–489.

Segal, I. (2003): "Optimal Pricing Mechanisms with Unknown Demand," *American Economic Review*, 93(3), 509–529.

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

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

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

B Online Appendix

B.1 Alternative estimators

We provide results using an alternative estimation technique based on the kernel density estimator.

An alternative nonparametric approach is to estimate cdfs using the empirical distributions defined previously, $\hat{F}(j) = \frac{n-j+1}{n+1}$ and $\hat{G}(j) = \frac{j}{m+1}$, and the densities using the kernel density estimator. For the kernel density estimator, we follow Guerre, Perrigne, and Vuong (2000) in using the triweight kernel, $\bar{k}(x) \equiv \frac{35}{32}(1-x^2)^3$, for $-1 < x < 1$, with bandwidth $\bar{h}(\mathbf{x}) \equiv 1.06\sigma_{\mathbf{x}}|\mathbf{x}|^{-\frac{1}{5}}$, where $\sigma_{\mathbf{x}}$ is the standard deviation of data vector $\mathbf{x}$. Thus, the unbounded kernel density estimator given data $\mathbf{x}$ is $\kappa_{\mathbf{x}}(y) \equiv \frac{1}{\bar{h}(\mathbf{x})|\mathbf{x}|} \sum_{i=1}^{|\mathbf{x}|} \bar{k}\left(\frac{y-x_{(i)}}{\bar{h}(\mathbf{x})}\right)$ and the kernel density estimator given data $\mathbf{x}$ and upper bound $\max_i x_i$ is calculated by reflecting the data across the upper bound and then truncating the resulting density, i.e., using $\kappa_{\mathbf{x}}^{ub}(y) \equiv 2\bar{\kappa}_{(\mathbf{x},-\mathbf{x}+2\max_i x_i))}(y)$, for $y \leq \max_i x_i$, and analogously for a lower bound, $\kappa_{\mathbf{x}}^{lb}(y) \equiv 2\bar{\kappa}_{(-\mathbf{x}+2\min_i x_i,\mathbf{x})}(y)$, for $y \geq \min_i x_i$.

Retaining our focus on privacy preserving implementation, we use observed types $\mathbf{v}_{(j)} \equiv (v_{(j)}, ..., v_{(n)})$ and $\mathbf{c}_{[j]} \equiv (c_{[j]}, ..., c_{[m]})$ to estimate the truncated density by imposing an upper bound $v_{(j)}$ for values and a lower bound $c_{[j]}$ for costs, and reweighting by $\hat{F}(j)$ and $1 - \hat{G}(j)$, respectively. Thus, we have:

$$\bar{f}(\mathbf{v}_{(j)}) = \hat{F}(n - |\mathbf{v}_{(j)}| + 1)\kappa_{\mathbf{v}_{(j)}}^{ub}(v_{(j)})$$

and

$$\bar{g}(\mathbf{c}_{[j]}) = (1 - \hat{G}(m - |\mathbf{c}_{[j]}| + 1))\kappa_{\mathbf{c}_{[j]}}^{lb}(c_{[j]}).$$

As an illustration, for $n = 20$ and values drawn from $U[0, 1]$, the estimated densities associated with various subsets of the data are shown in Figure 5.

We define the kernel-based weighted virtual types as

$$\bar{\Phi}_\alpha(\mathbf{v}_{(j)}) \equiv \begin{cases} v_{(j)} - \alpha\frac{1-\hat{F}(n-|\mathbf{v}_{(j)}|+1)}{\bar{f}(\mathbf{v}_{(j)})}, & \text{if } j < n \\ v_{(j)} - \alpha\frac{n-2}{n+1}, & \text{otherwise} \end{cases}$$

and

$$\bar{\Gamma}_\alpha(\mathbf{c}_{[j]}) \equiv \begin{cases} c_{[j]} + \alpha\frac{\hat{G}(m-|\mathbf{c}_{[j]}|+1)}{\bar{g}(\mathbf{c}_{[j]})}, & \text{if } j < m \\ c_{[j]} + \alpha\frac{m-2}{m+1}, & \text{otherwise,} \end{cases}$$

where for the case of $j = n$ or $j = m$, the kernel density estimator is undefined, so we use the same estimator as in our augmented prior-free mechanism.

The kernel-based double-clock auction is defined by comparison and target functions

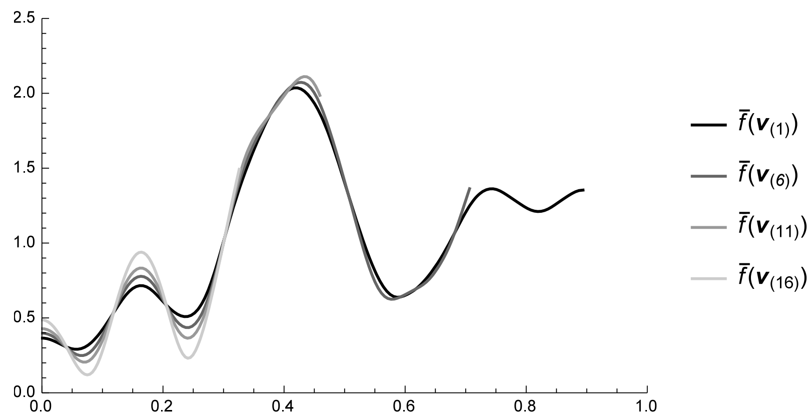

Figure 5: Kernel-based estimates of density f based on $\mathbf{v}_{(j)}$ for various j with $n = 20$ and values drawn from $U[0, 1]$.

given by, for $\boldsymbol{\omega}_t^B = (\mathbb{N}^A, \mathbf{x}^B, p^B)$ and $\boldsymbol{\omega}_t^S = (\mathbb{M}^A, \mathbf{x}^S, p^S)$, with $p^B \geq \max_i x_i^B$ and $p^S \leq \min_i x_i^S$, and $n^A = m^A$,

$$\phi(p^B, \boldsymbol{\omega}_t) = \bar{\Phi}_\alpha(p^B, x_{(2)}^B, ..., x_{(n)}^B) \quad \text{and} \quad \gamma(p^S, \boldsymbol{\omega}_t) = \bar{\Gamma}_\alpha(p^S, x_{[2]}^S, ..., x_{[m]}^S),$$

and

$$\tau(\boldsymbol{\omega}_t) = \bar{\theta}_{n^A},$$

where for $j < \min\{m, n\}$, we define $\bar{\theta}_j$ such that

$$\frac{\bar{\theta}_j - \bar{\Phi}_\alpha(\mathbf{v}_{(j+1)})}{\frac{1-\hat{F}(n-|\mathbf{v}_{(j+1)}|+1)}{\bar{f}(\mathbf{v}_{(j+1)})}} = \frac{\bar{\Gamma}_\alpha(\mathbf{c}_{[j+1]}) - \bar{\theta}_j}{\frac{\hat{G}(m-|\mathbf{c}_{[j+1]}|+1)}{\bar{g}(\mathbf{c}_{[j+1]})}}, \tag{12}$$

if such a $\bar{\theta}_j \in [\bar{\Phi}_\alpha(\mathbf{v}_{(j+1)}), \bar{\Gamma}_\alpha(\mathbf{c}_{[j+1]})]$ exists, and otherwise if the left side of (12) is greater than the right for all $\theta \in [\bar{\Phi}_\alpha(\mathbf{v}_{(j+1)}), \bar{\Gamma}_\alpha(\mathbf{c}_{[j+1]})]$, then $\bar{\theta}_j = \bar{\Gamma}_\alpha(\mathbf{c}_{[j+1]})$, and otherwise $\bar{\theta}_j = \bar{\Phi}_\alpha(\mathbf{v}_{(j+1)})$.

Figure 6(a) shows that the kernel-based virtual types are similar to the smoothed virtual types, although for the example shown they have greater volatility for high indices when the estimates are based on a small number of observations.

The fact that volatility decreases as the index decreases in Figure 6(a) suggests that our result on uniform convergence in probability would go through with the kernel-based virtual types, and so the asymptotic optimality results would go through as well. This is borne out in Figure 6(b), which illustrates the performance of the kernel-based mechanism relative to the BOPP mechanism and compared with the augmented prior-free mechanism. As shown in Figure 6(b), the kernel-based mechanism delivers performance similar

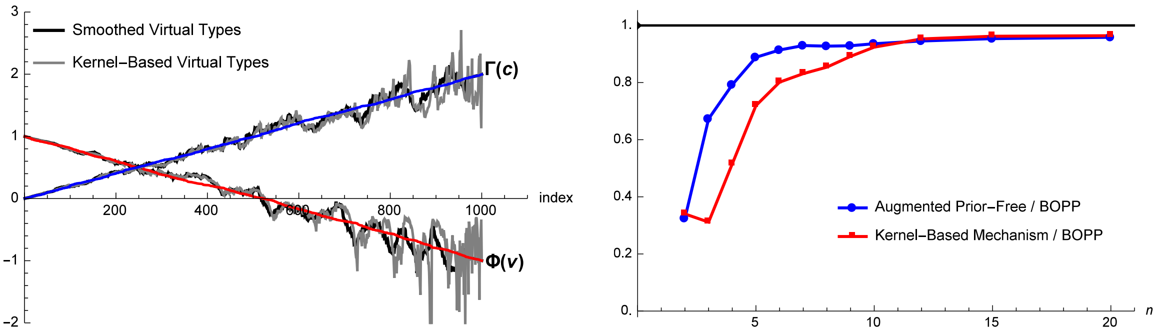

(a) Comparison of virtual types

(b) Comparison of mechanisms

Figure 6: Panel (a): Theoretical virtual types Φ and Γ, smoothed virtual types, and kernel-based virtual types with $\alpha = 1$ for 1000 values and 1000 costs drawn from the Uniform distribution on $[0, 1]$. Panel (b): Ratio of augmented prior-free and kernel-based expected revenue to BOPP expected revenue for various numbers of symmetric agents based on Monte Carlo simulation (5000 auctions) with buyers' values and sellers' costs drawn from the Uniform distribution on $[0, 1]$.

to the augmented prior-free mechanism when there are 10 or more buyer-seller pairs, but the kernel-based mechanism does not perform as well as the augmented prior-free mechanism for smaller numbers of agents (except for $n = 2$ where the mechanisms are defined similarly). This is consistent with the higher volatility for the kernel-based virtual types, which result in a higher that optimal number of trades in the kernel-based mechanism.

B.2 Rate of convergence

We approach an analysis of the rate of convergence along the lines of Satterthwaite and Williams (1989) and McAfee (1992).

For the case of $\alpha = 0$, the results of McAfee (1992) imply that under conditions that imply that spacings between order statistics are on the order of $1/\min\{m, n\}$, the expected efficiency loss in the baseline prior-free mechanism is on the order $1/\min\{m, n\}$, and the expected efficiency loss per potential trader is on the order $1/(\min\{m, n\})^2$.[31]

When $\alpha = 0$, the number of trades in the baseline prior-free mechanism differs from the number in the efficient mechanism by at most one trade. Focusing on revenue ($\alpha = 1$), for large numbers of traders, with high probability the number of trades in the baseline prior-free mechanism also differs from the number in the optimal mechanism by at most

[31] Rustichini, Satterthwaite, and Williams (1994) show that if $F = G$, then the market maker's profit in McAfee's mechanism is bounded, and hence has no effect in the large economy on surplus (i.e., does not affect the rate of convergence). Because we include revenue as part of the designer's objective, we include the market maker's gain when evaluating the asymptotic properties of a mechanism.

one trade. If k is the optimal number of trades, then the revenue impact from having one too many trades is that buyers pay $v_{(k+2)}$ instead of $v_{(k+1)}$ and sellers receive $c_{[k+2]}$ instead of $c_{[k+1]}$. The revenue loss is then

$$k\left(v_{(k+1)} - c_{[k+1]}\right) - (k+1)\left(v_{(k+2)} - c_{[k+2]}\right),$$

which for large numbers of traders is approximately $k\left(v_{(k+1)} - c_{[k+1]}\right)$, suggesting that the revenue loss per potential trader is on the order of $1/\min\{m, n\}$. This rate of convergence is borne out for the distributions we have considered. For example, as shown in Figure 7 for four distributional assumptions, the gap between limiting per-buyer revenue in the optimal mechanism and expected per-buyer revenue in the baseline prior-free mechanism for a given number of buyer-seller pairs n remains below $1/n$ for n up to 1000.

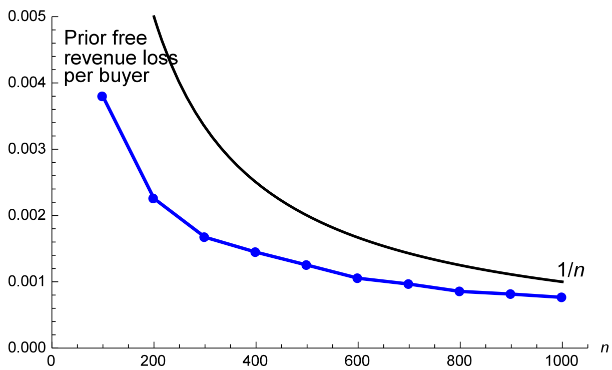

(a) Buyers and sellers: U(0,1)

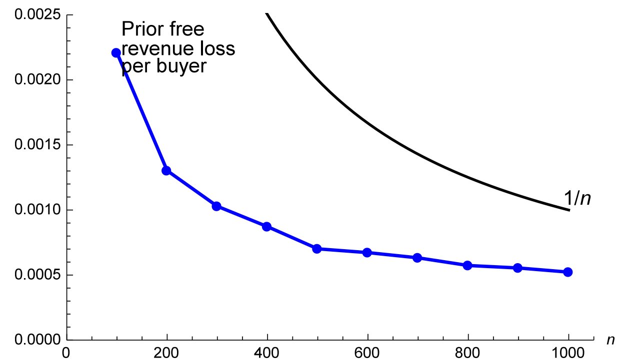

(b) Buyers and sellers: Beta(2,4)

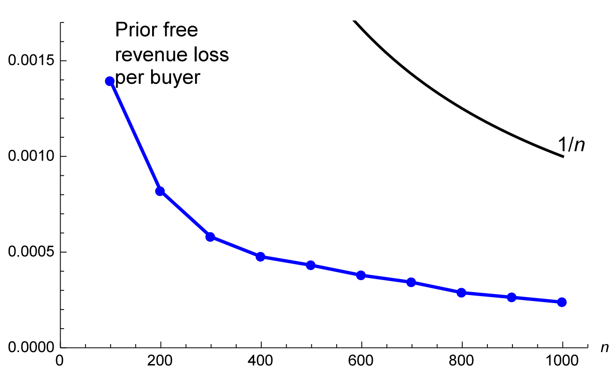

(c) Buyers: Beta(2,4), Sellers: Beta(4,2)

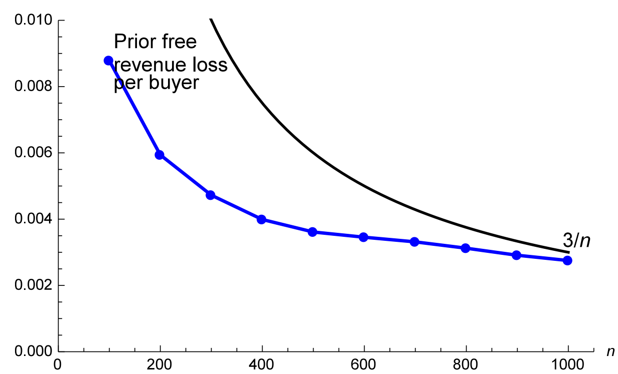

(d) Buyers and sellers: Exponential(1)

Figure 7: Revenue loss per buyer in the baseline prior-free mechanism relative to the limiting optimal revenue per buyer (0.125 for panel (a), 0.065 for panel (b), 0.018 for panel (c), and 0.278 for panel (d)) for various numbers of symmetric agents based on Monte Carlo simulation (5000 auctions) with buyers' values and sellers' costs drawn from the distributions indicated.

B.3 Value of estimation when $\alpha = 0$

We now briefly illustrate that estimation can improve outcomes because it allows the target function to vary with observable details. As a case in point, consider the mechanism of McAfee (1992), which does not rely on estimation and uses a target function with $(v_{(j)} + c_{[j]})/2$, thereby taking the midpoint of the interval between the last exit prices as the target. Although this target is a good approximation if values and costs are similarly spaced, for example because the numbers of buyers and sellers and their distributions are similar, it lacks the flexibility to vary with the details of the environment otherwise.

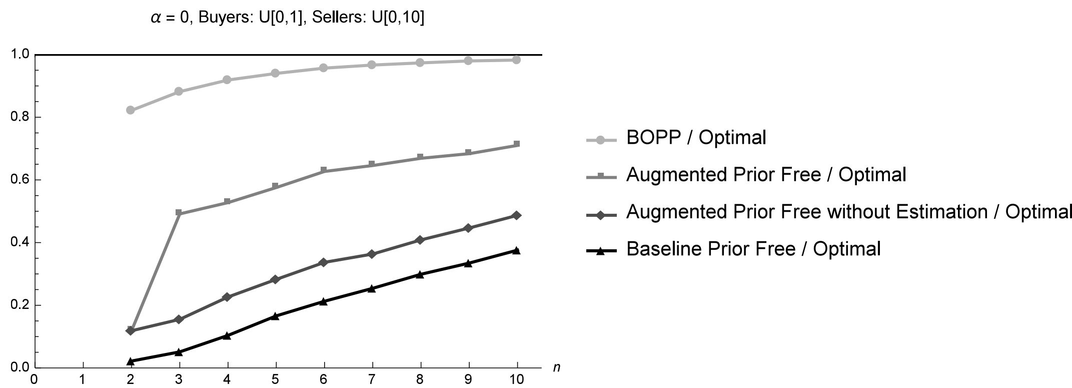

Figure 8: Illustration of the value of estimation for $\alpha = 0$.

To illustrate this, assume $n = m$ and that buyers draw their values from the Uniform distribution on $[0, 1]$ whereas sellers draw their costs from the Uniform distribution on $[0, 10]$. In this case, for a given number of active buyer-seller pairs, buyers' values are in expectation spaced more closely together than sellers' costs. The gap between the line "Augmented Prior Free/Optimal" and the line "Augmented Prior Free without Estimation/Optimal" in Figure 8 shows that estimation matters even for $\alpha = 0$. For example, for $n = 3$ and $n = 4$, the augmented prior-free mechanism with estimation achieves more than twice the social surplus of the augmented prior-free mechanism without estimation. For $n = 10$, the mechanism with estimation still outperforms the one without by 43%.

The mechanism that implements the target functions without estimation by using the midpoint of the last observed exits performs poorly because in many cases the midpoint will be greater than the upper support of the buyers' distribution. This means that the target cannot be achieved. Using estimation, the observed closer spacings between values than between costs (for $n > 2$) results in a target that is relatively closer to (and perhaps even equal to) the last observed value, and so more likely to be in a range that can be achieved. This explains the better performance (for $n > 2$) of the augmented prior-free mechanism relative to the corresponding mechanism without estimation.