Games and Economic Behavior 137 (2023) 68–90

Contents lists available at ScienceDirect

Games and Economic Behavior

journal homepage: www.elsevier.com/locate/geb

Asymptotically optimal prior-free asset market mechanisms ☆

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

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

ARTICLE INFO

Article history:
Received 1 September 2021
Available online 14 November 2022

JEL classification:
C72
D44
L13

Keywords:
Market mechanism
Endogenous trading position
Detail free
Mechanism design with estimation
Laffer curve

ABSTRACT

We develop a prior-free mechanism for an asset market that is dominant-strategy incentive compatible, ex post individually rational, constrained efficient, and asymptotically optimal—as the number of agents grows large, the designer's profit from using this mechanism approaches the profit it would optimally make if it knew the agents' type distribution at the outset. The direct implementation first identifies the agent whose value equals the Walrasian price. The second step can be described algorithmically as consisting of ascending and descending clock auctions that start from the Walrasian price, estimate virtual types, and stop eliminating trades when the estimated virtual value exceeds the estimated virtual cost. The mechanism permits partial clock auction implementation. Our approach accommodates heterogeneity among groups of traders and discrimination among these, provided heterogeneity is not too accentuated.

1. Introduction

Exchanges in which agents' trading positions—buy, sell, or hold—are determined endogenously abound in the real world. If anything, they have become even more prevalent in the digital age as companies like Uber, Lyft, and AirBnB allow agents who used to be active on only one side of the market as buyers or sellers to choose whether they consume or provide services. Moreover, vertical integration almost inevitably makes an integrated firm's trading position endogenous because, for some realizations of values, the vertically integrated firm may optimally buy additional inputs, sell its units, or not trade in the input market. The mechanism design literature has long paid scant attention to such markets, which we refer to as *asset markets*, but there has been a recent upsurge of interest.[1] An important issue, of particular salience in novel industries and environments, concerns the mechanism for a designer that faces traders without knowing the distributions of their types.

In this paper, we approach this problem by looking for mechanisms that satisfy dominant strategy incentive compatibility and individual rationality ex post, which seem the natural constraints in environments without knowledge of the

☆ We thank an Advisory Editor and a reviewer of this journal for comments and suggestions that have helped us to improve the paper and its exposition substantively. We are also grateful for valuable discussions with Julien Combe, Bing Liu, Toan Le, and Ellen Muir, and we thank seminar participants at the University of Melbourne. Financial support from the Samuel and June Hordern Endowment and the Australian Research Council under a Discovery Project Grant DP200103574 is gratefully acknowledged.

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

[1] See, for example, Lu and Robert (2001), Chen and Li (2018), Loertscher and Wasser (2019), Johnson (2019), Loertscher and Marx (2020b), Delacrétaz et al. (2022), and Li and Dworczak (2021), who build on variants of the partnership model of Cramton et al. (1987). The labels used to describe asset markets may differ from one paper to another.

https://doi.org/10.1016/j.geb.2022.10.013
0899-8256/

prior distribution, and achieve maximum revenue as the number of ex ante symmetric agents becomes large. That is, as the number of agents becomes large the mechanism converges to the one that maximizes the designer's profit if the designer knew the distribution from which the agents draw their types. We do this under the assumptions that the agents' endowments and maximum demands are commonly known, their constant marginal values are independent draws from the same distribution and that this (unknown) distribution is regular in the sense of exhibiting both increasing virtual costs and virtual values.

Moreover, we show that this *asymptotically optimal* prior-free mechanism, in addition, always satisfies *constrained efficiency* insofar as the agents who sell their endowments under this mechanism are the ones with the lowest values, and the agents who buy under this mechanism are the ones with the highest values. Constrained efficiency is a desirable property because it eliminates post-allocation gains from trade on each side of the market among agents that trade in the mechanism. Thereby, it reduces scope for resale and, related, bid shading and public fall-out due to discrimination that arises when, say, an agent who submitted a lower bid obtains a unit while another with a higher bid is not served (Loertscher and Mezzetti, 2021). In light of the Wilson critique (Wilson, 1987), being prior-free is a desirable robustness property, while asymptotic optimality insures the designer against ex post regret—the designer would not have made more profit, ex post, had it known at the outset the type distribution that it learned by running the mechanism (Loertscher and Marx, 2020a). The asymptotic optimality of a mechanism that always endows agents with dominant strategies and respects their individual rationality constraint is a property that is similar to the consistency of an estimator in econometrics that may also have other desirable features, such as unbiasedness. These properties can be established analytically, but are typically not interpreted as meaning that the estimator or mechanism only performs well in the large. Indeed, the simulations in Section 4 show that our prior-free mechanism performs well even with a moderate number of agents, and by construction, it always satisfies dominant strategy incentive compatibility, ex post individual rationality, and constrained efficiency.

A particular feature of asset markets and partnership models à la Cramton et al. (1987), and a major challenge to designing suitable mechanisms for these environments, is that they are subject to what Lewis and Sappington (1989) dubbed "countervailing incentives": agents who are likely to trade as sellers have incentives to exaggerate their costs, while agents with high values are likely to trade as buyers and therefore have incentives to report lower values than they have. This implies that in an asset market, an agent's information rent is typically minimized for a type in the interior of the type space. This contrasts sharply with one-sided allocation problems such as auctions or procurement auctions, as well as the two-sided allocation problems addressed by double auctions, where it is a priori known which agents are buyers and sellers and therefore a priori known which type of an agent is the worst-off—the lowest possible type for a buyer and the highest possible type for a seller.

From the designer's perspective, knowledge of which types are worst-off identifies the types for which the individual rationality constraint binds, which, in turn, permits clock auction implementation of the optimal mechanism in environments with one-sided and two-sided private information. Because of the countervailing incentives and the endogeneity of trading positions inherent in asset markets, these markets afford no simple, a priori known point at which one can, say, start a clock auction. The mechanism we develop is, therefore, direct. It first identifies agents with values equal to the Walrasian price defined with respect to the reported types. As the number of agents goes to infinity, these agents will almost surely not trade under the mechanism that is optimal for the market maker in the same sense and under the same conditions as Myerson's optimal auction is optimal for the auctioneer. Its second step can be described algorithmically as consisting of ascending and descending clock auctions that start from the Walrasian price and estimate virtual types along the way using the methodology developed in Loertscher and Marx (2020a).

Although the trades that are executed are efficient in that they involve sales from the lowest-type agents to the highest-type agents, the mechanism that maximizes the market maker's expected profit does not, in general, involve the efficient number of trades. Just as a monopolist sells less than the efficient quantity, our market maker generally restricts the quantity of trades below the efficient level. In contrast to an optimal sales (procurement) auction, the distortion away from the efficient allocation does not involve traders with low (high) values, but rather agents of intermediate types equal to or near the Walrasian price. These are the types for whom the individual rationality constraint is binding in an optimal mechanism.

As we show, our mechanism permits partial clock implementation, which is desirable because, as argued by Baranov et al. (2017), even partial privacy preservation is valuable when there is a tension between privacy preservation and price discovery. For example, one can start an ascending clock. The clock increases as long as indicated net demand exceeds supply. An agent that exits either does not trade or trades as a seller at a price that does not depend on its bid (other than its having exited). As soon as net demand is less than supply, the mechanism begins estimating virtual values and virtual costs, beginning in both directions from the price at which net demand equals supply. The clock auction continues to reduce the quantity demanded until the first point at which estimated virtual values exceed estimated virtual costs. This ascending clock auction preserves the privacy of agents who trade as buyers and endows agents with obviously dominant strategies in the subgame that ensues after the Walrasian price is reached. Conversely, one can implement the mechanism using a descending clock auction and preserve the privacy of the agents who trade as sellers. Of course, if the Walrasian price for the problem at hand is known a priori, for example based on historical data, then there is an asymptotically optimal mechanism that requires no prior information about distributions and that can be implemented using clock auctions by starting both the buyers' and the sellers' clock auctions at that Walrasian price.

The main analysis assumes that all agents draw their types from the same distribution and have the same endowments and maximum demands, where for simplicity we set the ratio of maximum demand over endowment equal to two. In an extension, we show that the assumptions on maximum demands and endowments can be relaxed straightforwardly. We also provide conditions under which one can allow for heterogeneous groups of traders, with traders in each group drawing their types from the same regular distribution, while the distributions are allowed to vary across groups. Last but not least, we show that the problem faced by the designer who aims to maximize revenue in the large is equivalent to the problem faced by a tax-revenue maximizing authority that does not know the shape of the Laffer curve (Laffer, 2004), and we explain how the mechanism can be augmented to accommodate a Ramsey objective consisting of a convex combination of revenue and social surplus.

To the best of our knowledge, the mechanism design literature on asset markets is confined to Lu and Robert (2001), Chen and Li (2018), Loertscher and Wasser (2019), Johnson (2019), Loertscher and Marx (2020b), Delacrétaz et al. (2022), and Li and Dworczak (2021), which, as mentioned in footnote 1, build on Cramton et al. (1987). None of these analyzes prior-free asymptotic optimality.[2] While details differ (in particular, there are no technological interference constraints and agents do not register as buyers or sellers here), this mechanism is reminiscent of the incentive auction employed by the U.S. Federal Communications Commission in 2017 to purchase spectrum licenses from television broadcasters and then sell them to mobile wireless providers (see, e.g., Leyton-Brown et al., 2017; Milgrom, 2017).

Outside the domain of asset markets and beginning with Baliga and Vohra (2003) and Segal (2003), there is a sizeable literature on prior-free asymptotically optimal mechanisms.[3] To maintain incentive compatibility when agents compete against each other for the opportunity to trade, most of this literature takes either a random-sampling approach, whereby a small subset of agents is sampled and prevented from trading, or splits the market into separate submarkets and uses reports and estimates from one submarket to determine the allocation and payments in the other. Neither of these mechanisms is constrained efficient. An exception is Loertscher and Marx (2020a), which develops a clock auction that is prior free and asymptotically optimal if the type distributions satisfy regularity. In this paper, we use the estimation and elimination procedure developed in Loertscher and Marx (2020a) in the second stage of a direct mechanism, which in a first stage needs to identify the Walrasian price in order to initiate that procedure. The first stage of this mechanism embeds elements from the trade-sacrifice mechanism that we developed in Loertscher and Marx (2020b), but the problems addressed by that paper and the present one differ substantively. In particular, agents with types equal or close to the Walrasian price never trade in the large-market limit here, while (some) agents with types close to the Walrasian price always trade under efficiency (and in the trade sacrifice mechanism) in Loertscher and Marx (2020b). A key challenge for making things work for asymptotic optimality in the asset market setting is that there is no obvious starting point for the mechanism. While a designer selling to multiple buyers can use an ascending clock auction that starts at a very low price, and a designer purchasing from multiple suppliers can use an descending clock auction that starts at a very high price, in the asset market setting it is not clear ex ante which agents should be buyers and which should be suppliers (and which agents should not trade).

More broadly and fundamentally, the paper contributes to what may be a nascent field in policy making and market design that recognizes the benefits of adjusting allocations in real time to information gleaned from earlier allocations. Chetty et al. (2021) provide a simple yet powerful illustration of this in the context of pandemic-related stimulus checks and present persuasive arguments for why reliance on past experience is neither useful nor necessary.[4] A natural reinterpretation of our model and application to public finance is to view it as a Laffer curve problem, where the designer—that is, the tax authority—needs to find the tax rate that maximizes revenue, which in an evolving environment requires estimating the elasticity of the volume traded with respect to the tax rate. In a different context, Agarwal et al. (2021) show how waitlists whose priorities depend on empirical distributions can improve outcomes for patients waiting for a cadaver kidney.[5]

The remainder of this paper is organized as follows. Section 2 introduces the setup. In Section 3, we derive asymptotically optimal, prior-free mechanisms. Section 4 analyzes performance in the small. Section 5 contains extensions and discussion, and Section 6 concludes the paper.

2. Setup

We assume that there is a set $\mathcal{N}$ of $n \geq 3$ agents, where n is odd. Each agent $i \in \mathcal{N}$ has an endowment of one unit of a homogeneous good and a constant marginal value θ_i for up to two units of the good. The willingness to pay for additional units beyond the capacity is zero. The agents' endowments and maximum demands are commonly known, including by the market maker (or designer). In this setup, efficiency would be achieved by having the $\frac{n-1}{2}$ agents with the lowest values for the good transfer their units to the $\frac{n-1}{2}$ agents with the highest values for the good. But we consider a market maker that

[2] Interestingly, the idea of estimating the agents' type distribution was mentioned by Cramton et al. (1987, p. 624) without being further pursued there.

[3] See, for example, Loertscher and Marx (2020a) for a comprehensive list of references.

[4] Chetty et al. (2021) emphasize the value of real-time data on how consumers have spent stimulus checks for the design of continuing stimulus, concluding that "targeting stimulus checks to lower- and middle-income households and using the money that is saved for programs to support those who need the most help is likely to yield greater economic benefits."

[5] The recent work by Abdulkadiroğlu et al. (2017, 2022), who use tie-breaking in centralized matching mechanisms for identification and estimation, is also related to our paper. One difference is that in our mechanism the estimates feed back into the allocation and transfers.

is not interested in efficiency, but rather in maximizing its own profit.[6] Further, we assume that the agents' values for the good, which we refer to as the agents' types, are their own private information, so incentives must be provided for agents to reveal the relevant information, and consequently, information rents will need to be paid.

It will be useful to have notation for order statistics among the agents' types. For $k \in \{1, \dots, \frac{n-1}{2}\}$, we use $\theta_{(k)}$ to denote the k-th highest type and $\theta_{[k]}$ to denote the k-th lowest type in $\boldsymbol{\theta} \equiv (\theta_1, \dots, \theta_n)$. We define $\theta_{(0)} = 1 = \theta_{[\frac{n+1}{2}]}$ and $\theta_{(\frac{n+1}{2})} = 0 = \theta_{[0]}$.

We assume that the agents' types are drawn independently from a continuously differentiable distribution F with compact interval support, which we normalize without loss of generality to $[0,1]$, and density f that is positive on $[0,1]$. To simplify notation, we assume that there are no ties among the agents' types.[7] We assume further that F has increasing virtual type functions

$$\Phi(\theta) \equiv \theta - \frac{1-F(\theta)}{f(\theta)} \text{ and } \Gamma(\theta) \equiv \theta + \frac{F(\theta)}{f(\theta)}, \tag{1}$$

where we refer to Φ as the virtual value function and to Γ as the virtual cost function. As is well known and will become clear from the analysis below, Φ can be interpreted as a marginal revenue function and Γ as a marginal cost of procurement function.[8]

We assume that neither the market maker nor the agents know F. That the agents do not know F is in the spirit of robust mechanism design, and we impose it to prevent the market maker from simply asking agents to report F and if they fail to agree, to prohibit all trade. However, we assume that the market maker knows that F has a continuous, positive density on $[0,1]$ and increasing virtual type functions.

We say a mechanism is *prior-free* if it is defined without reference to the type distribution. A direct mechanism $\langle \mathbf{s}, \mathbf{p} \rangle$ consists of an allocation rule $\mathbf{s} : [0,1]^n \to \{0,1,2\}^n$ satisfying $\sum_{i \in \mathcal{N}} s_i(\boldsymbol{\theta}) = n$ for all $\boldsymbol{\theta} \in [0,1]^n$ and a payment rule $\mathbf{p} : [0,1]^n \to \mathbb{R}^n$. The mechanism satisfies dominant-strategy incentive compatibility (DIC) if, for all $i \in \mathcal{N}$, all $\theta_i, \hat{\theta}_i \in [0,1]$, and all $\boldsymbol{\theta}_{-i} \in [0,1]^{n-1}$,

$$s_i(\boldsymbol{\theta})\theta_i - p_i(\boldsymbol{\theta}) \geq s_i(\hat{\theta}_i, \boldsymbol{\theta}_{-i})\theta_i - p_i(\hat{\theta}_i, \boldsymbol{\theta}_{-i}),$$

where $\boldsymbol{\theta} = (\theta_i, \boldsymbol{\theta}_{-i})$, and ex post individual rationality (EIR) if for all $i \in \mathcal{N}$, all $\theta_i \in [0,1]$, and all $\boldsymbol{\theta}_{-i} \in [0,1]^{n-1}$,

$$s_i(\boldsymbol{\theta})\theta_i - p_i(\boldsymbol{\theta}) \geq \theta_i,$$

where θ_i is the value of the outside option of agent i, which is the value of consuming its endowment.

If the market maker knew the distribution F and faced a continuum of traders with mass normalized to 1, its profit maximization problem would be to choose the quantity $Q \in [0,1]$ to maximize

$$(F^{-1}(1-Q) - F^{-1}(Q))Q,$$

where $F^{-1}(1-Q)$ is such that a mass $1-Q$ of agents have higher values than $F^{-1}(1-Q)$ and $F^{-1}(Q)$ is such that a mass Q of agents have values that are smaller than $F^{-1}(Q)$. Setting $p^B = F^{-1}(1-Q)$ and $p^S = F^{-1}(Q)$, the first-order condition for a profit maximum is

$$\Phi(p^B) = \Gamma(p^S), \tag{2}$$

and $p^B = F^{-1}(1-Q^*)$ and $p^S = F^{-1}(Q^*)$ is equivalent to the market-clearing condition

$$1 - F(p^B) = Q^* = F(p^S). \tag{3}$$

It is straightforward to verify that the second-order condition is satisfied whenever the first-order condition is satisfied if Φ and Γ are both monotone, which establishes that (2) and (3) characterize the unique maximum. Intuitively, (2) equalizes the marginal revenue with the marginal cost of procurement and thereby illustrates the aforementioned interpretations of Φ and Γ. It follows that per-trader profit in the large-market limit when the market maker knows F is $(p^B - p^S)Q^*$.[9] Fig. 1 provides an illustration for the uniform distribution. We refer to mechanisms that maximize the market maker's expected profit, with the expectation being taken with respect to the true distribution F, as *optimal* mechanisms.

We say that a DIC-EIR mechanism is *asymptotically optimal* if its per-trader profit converges in probability to $(p^B - p^S)Q^*$ as n (with n odd) goes to infinity. As mentioned, the property of asymptotic optimality provides valuable discipline for a

[6] It is straightforward to allow the market maker's objective to be the maximization of the expected value of a weighted average of profit and social surplus, as described in Section 5.1.

[7] One can accommodate ties by introducing an ordering over agents and using that ordering to define the tie-breaking rule; see, for example, Loertscher and Marx (2020b).

[8] See, for example, Mussa and Rosen (1978) and Bulow and Roberts (1989).

[9] See Corollary B.1 in Appendix B, which extends Lu and Robert (2001, Theorem 3) to the limit case for our setup.

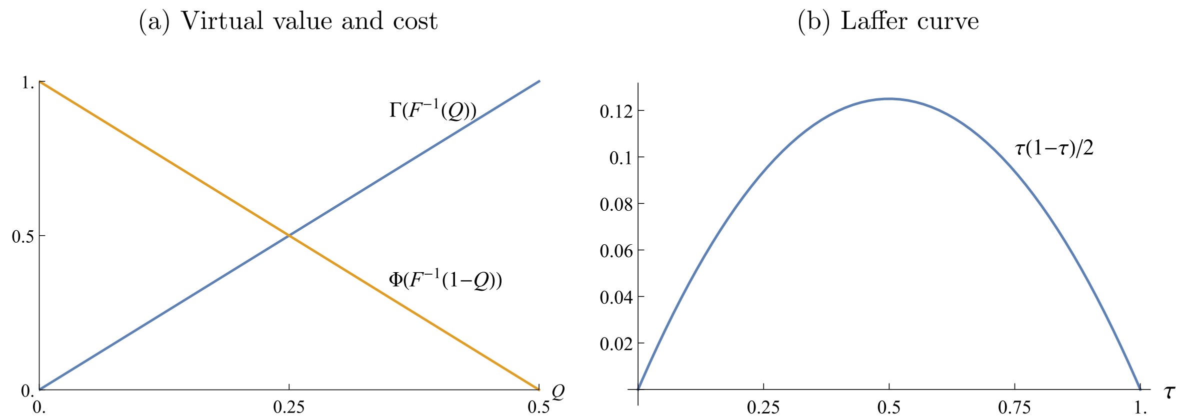

Fig. 1. Illustrations for uniformly distributed types. With uniform types, $\Phi(\theta)=2\theta-1$ and $\Gamma(\theta)=2\theta$, and so $\Phi(F^{-1}(1-Q))=1-2Q$ and $\Gamma(F^{-1}(Q))=2Q$. Thus $Q^*=1/4$, $p^B=3/4$, $p^S=1/4$, and $(p^B-p^S)Q^*=1/8$ as illustrated in panel (a). Panel (b) depicts the equivalent representation as a Laffer curve problem, where τ is as specific tax and tax revenue is $\tau(1-\tau)/2$.

mechanism, like requiring consistency for a statistic being used as an estimator. Asymptotic optimality protects the market maker from ex post regret because it means that the market maker could not have done better had it known from the outset what it learned about distributions and virtual types from running the mechanism.

When the market maker does not know F, but knows that it exhibits monotone virtual type functions, as we assume, its problem boils down to estimating when the point of intersection between $\Phi(F^{-1}(1-Q))$ and $\Gamma(F^{-1}(Q))$ has been reached. Asymptotic optimality then requires this estimation to be accurate as the number of traders goes to infinity. Of course, the problem of empirically finding the maximum of a function is reminiscent of the discussions surrounding the Laffer curve (Laffer, 2004), which gives tax revenue as a function of the tax rate. Indeed, if one imposes a specific tax $\tau \geq 0$, the tax revenue with a continuum of traders drawing their types from the uniform distribution on $[0,1]$ is $\tau(1-\tau)/2$ as illustrated in panel (b) of Fig. 1.[10] We return to the Laffer curve problem—that is, the problem of identifying its top—in Section 5.1.

This setup provides a useful starting point for the analysis. Profit maximization comes down to the selection of which agents should sell their endowments and which agents should buy those units. While the underlying economics are most transparent in this simple setup, the assumption that n is odd simplifies the analysis without being essential, as discussed at the end of Section 3.1, and the assumptions that endowments and maximal demands are the same across all agents (and equal to one and two, respectively) are also not essential, as we show in Section 5.4.

3. Asymptotically optimal prior-free mechanisms

We begin by defining a prior-free, incentive compatible, individually rational asset market mechanism with estimated virtual types, and we then define the prior-free estimators for the virtual types to be used in that mechanism. Of course, if the mechanism relies only on *estimated* virtual types, then it is prior free, and if those estimators are consistent, then we have a prior-free mechanism that is asymptotically optimal. We conclude the section by showing that the prior-free asset market mechanism with estimated virtual types can be implemented using a partial clock auction that preserves the privacy of trading agents on one side of the market. For practical implementation, privacy preservation can be desirable because it protects agents from the exploitation of their private information.

3.1. Prior-free mechanism with estimated virtual types

Although the prior-free asset market mechanism with estimated virtual types that we define is a direct mechanism that requires agents to report their types only once, at the outset, one can usefully think of it as consisting of two phases.

In the first phase, given reports $\theta_1,\ldots,\theta_n$, the mechanism determines the type of the agent that, under efficiency, does not trade: $\hat{\theta} \equiv \theta_{\left(\frac{n+1}{2}\right)}$. We refer to $\hat{\theta}$ as the *cutoff type* and the agent with type $\hat{\theta}$ as the *cutoff agent*. Note that $\hat{\theta}$ is the (unique) market clearing price given the realization $\boldsymbol{\theta}$. After the cutoff type is identified, the mechanism has a second phase, the "auction phase," in which agents with initial reports above the cutoff type participate as buyers in an ascending clock

[10] To see the equivalence between our problem and the Laffer problem, note that for the uniform distribution, $p^B-p^S=F^{-1}(1-Q^*)-F^{-1}(Q^*)=1-2Q^*$, and so the market maker's profit is $(p^B-p^S)Q^*=(1-2Q^*)Q^*$. Using the change of variables, $\tau=2Q^*$, this is equal to $(1-\tau)\tau/2$. Giving this tax interpretation, a tax τ on trade, borne symmetrically by buyers and sellers, means that for uniformly distributed types, trade only occurs between buyers with types greater than $1/2+\tau/2$ and sellers with types less than $1/2-\tau/2$, in which case the gains from trade exceed τ. Thus, the tax τ is collected on $(1-\tau)/2$ trades, for total tax revenue of $\tau(1-\tau)/2$.

auction starting at the cutoff type, while agents with reports below the cutoff type participate as sellers in a descending clock auction. In this auction phase, for $t \in \{0, 1, \dots\}$,[11] we have a state $\boldsymbol{\omega}_t = (\mathcal{B}_t, \mathcal{S}_t, p_t^B, p_t^S)$, where the components of the state are: the set of active buyers $\mathcal{B}_t \subseteq \mathcal{N}$, the set of active sellers $\mathcal{S}_t \subseteq \mathcal{N}$, which satisfy $\mathcal{B}_t \cap \mathcal{S}_t = \emptyset$, the buyer price $p_t^B \in [0, 1]$, and the seller price $p_t^S \in [0, 1]$. Letting Ω be the set of all such states, the auction phase relies on a virtual value estimator $\phi : \Omega \to \mathbb{R}$ that is increasing in the buyer price and a virtual cost estimator $\gamma : \Omega \to \mathbb{R}$ that is increasing in the seller price. We defer to Section 3.2 the task of defining consistent virtual type estimators.

We initialize the auction phase by setting $\mathcal{S}_0 \equiv \{i \in \mathcal{N} \mid \theta_i < \hat{\theta}\}$, $\mathcal{B}_0 \equiv \{i \in \mathcal{N} \mid \theta_i > \hat{\theta}\}$, and $p_0^B = p_0^S \equiv \hat{\theta}$. The cutoff agent is not included in either of the sets $\mathcal{S}_0$ or $\mathcal{B}_0$—as a result, the cutoff agent never trades. Also, $|\mathcal{B}_0| = |\mathcal{S}_0| = \frac{n-1}{2} \geq 1$, so, given that $n \geq 3$, the auction starts with at least one active buyer and seller.

For $t \in \{0, 1, \dots\}$, the auction phase applies the following iterative process:

- Step 1: If either $\phi(\boldsymbol{\omega}_t) \geq \gamma(\boldsymbol{\omega}_t)$ or $|\mathcal{B}_t| = 1$, then the auction phase ends; otherwise, proceed to step 2.
- Step 2: Let

$$\boldsymbol{\omega}_{t+1} = \left(\mathcal{B}_t \backslash \{j^B\},\ \mathcal{S}_t \backslash \{j^S\},\ \theta_{j^B},\ \theta_{j^S}\right),$$

  where j^B and j^S are defined by $\theta_{j^B} = \min_{i \in \mathcal{B}_t}\{\theta_i\}$ and $\theta_{j^S} = \max_{i \in \mathcal{S}_t}\{\theta_i\}$, and proceed to step 3.
- Step 3: Increment t by 1 and return to step 1.

Because step 2 reduces the number of active buyers by 1 and the number of active sellers by 1, the auction phase ends after a finite number of iterations. Let $\hat{\boldsymbol{\omega}} = (\hat{\mathcal{B}}, \hat{\mathcal{S}}, \hat{p}^B, \hat{p}^S)$ denote the state when the auction phase ends. By construction, $|\hat{\mathcal{B}}| = |\hat{\mathcal{S}}|$, $\hat{p}^B \geq \hat{\theta}$, and $\hat{p}^S \leq \hat{\theta}$. The asset market mechanism then specifies that $|\hat{\mathcal{B}}|$ units are traded, with each agent $i \in \hat{\mathcal{B}}$ receiving one unit and paying $\hat{p}^B$, and each agent $i \in \hat{\mathcal{S}}$ providing its endowment and being paid $\hat{p}^S$.

Proposition 1. *The prior-free asset market mechanism with virtual type estimators ϕ and γ is dominant-strategy incentive compatible, individually rational, and deficit free if for any state $\boldsymbol{\omega} = (\mathcal{B}, \mathcal{S}, p^B, p^S)$, $\phi(\boldsymbol{\omega})$ and $\gamma(\boldsymbol{\omega})$ depend only on p^B, p^S, and the types of inactive agents, that is of agents in $\mathcal{N} \backslash (\mathcal{B} \cup \mathcal{S})$.*

Proof of Proposition 1. If $\phi(\boldsymbol{\omega})$ and $\gamma(\boldsymbol{\omega})$ depend only on p^B and p^S, which are functions of the types of agents in $\mathcal{N} \backslash (\mathcal{B} \cup \mathcal{S})$, and on the types of agents in $\mathcal{N} \backslash (\mathcal{B} \cup \mathcal{S})$, then the stopping rule in the auction phase depends only on the types of agents that do not trade and so the auction phase is a two-sided clock auction as defined by Loertscher and Marx (2020a).[12] By the nature of clock auctions, it follows that in the auction phase every agent has a dominant strategy to bid truthfully. Hence, we are left to rule out profitable deviations from truthful reporting in the first phase, which determines whether an agent is the cutoff agent (and so does not trade) or participates as a buyer or a seller.

Consider first an agent who under truthful reporting is the cutoff agent. If the agent deviates from truthful reporting but remains the cutoff agent, its payoff remains zero and so the deviation is not profitable. If the agent reports a higher type so that it becomes a buyer, the cutoff type increases (above the deviating agent's type), so that if the agent trades in the clock auction, its payoff is negative. The best case for the agent is that it does not trade in the clock auction, in which case its payoff is zero. Thus, a deviation that causes the agent to become a buyer in the clock auction is not profitable. A symmetric argument shows that a deviation that causes the agent to become a seller in the clock auction is also not profitable.

We are left to consider agents who under truthful reporting are buyers or sellers in the clock auction. The smallest payoff from truthful reporting in the clock auction is zero, so the deviation to become the cutoff agent is not profitable. Any deviation that leaves an agent on the same side of the cutoff agent does not affect the cutoff type, and so has no effect on the starting prices of the clock auctions. Thus, no such deviation is profitable. The only deviations we are left to rule out as profitable are such that an agent who under truthful reporting is a buyer (seller) becomes a seller (buyer). But no such deviation can be profitable because the best that could happen to the deviating agent is that it does not trade, in which case its payoff is no more than the minimum payoff from truthful reporting in the clock auction. If the agent does trade, it either sells at a price below its true type or buys at a price above its true type, in which case it makes a loss.

Thus, the mechanism is dominant-strategy incentive compatible. That it is individually rational follows from the fact by bidding truthfully that every agent obtains a net payoff of at least zero. Deficit freeness follows from the facts that the lowest price paid by any buyer on any unit traded is the cutoff type, while the highest price received by any seller on any unit is the cutoff type, and that the quantity traded is balanced. ■

Proposition 1 puts us in a position to construct an asymptotically optimal prior-free market mechanism, provided that we have mappings ϕ and γ that give consistent estimates of the marginal buyer's virtual value and the marginal seller's virtual cost, respectively, without interfering with the incentive compatibility of the mechanism. The critical issue, then, is to construct consistent, prior-free estimates of the hazard rates $(1 - F(\theta))/f(\theta)$ and $F(\theta)/f(\theta)$ using only data from

[11] The mechanism that we define has no more that $\frac{n-3}{2}$ iterations, so t never exceeds $\frac{n-3}{2}$.

[12] The two-sided clock auction defined in Loertscher and Marx (2020a) is an adaptation of the one-sided clock auction of Milgrom and Segal (2020).

agents that do not trade. As we show below, the spacings-based estimator of Loertscher and Marx (2020a) provides such an estimator.

While the analysis above assumes that n is odd, one can easily accommodate an even number of agents n with $n \geq 4$. To do so, define the two "middle" types, $\hat{\theta}_1 \equiv \theta_{(n/2+1)}$ and $\hat{\theta}_2 \equiv \theta_{(n/2)}$, where $\hat{\theta}_1 \leq \hat{\theta}_2$, and initialize the auction phase with $p_0^S \equiv \hat{\theta}_1$, $p_0^B \equiv \hat{\theta}_2$, $\mathcal{S}_0 \equiv \{i \in \mathcal{N} \mid \theta_i < \hat{\theta}_1\}$, and $\mathcal{B}_0 \equiv \{i \in \mathcal{N} \mid \theta_i > \hat{\theta}_2\}$.

3.2. Asymptotic optimality

We now specify virtual type estimators ϕ_σ and γ_σ that satisfy the conditions of Proposition 1 and that have properties sufficient to guarantee asymptotic optimality.[13] Given a state $\boldsymbol{\omega} = (\mathcal{B}, \mathcal{S}, p^B, p^S)$, as one might expect, and analogous to how the analytic virtual value and cost functions are defined in (1), we can estimate the virtual value as p^B minus an estimate of a buyer's hazard rate and the virtual cost as p^S plus an estimate of the seller's hazard rate. Thus, the problem reduces to finding consistent estimates of the hazard rate terms in the virtual type functions.

Following Loertscher and Marx (2020a), for any state $\boldsymbol{\omega} = (\mathcal{B}, \mathcal{S}, p^B, p^S)$ with the number of active agents on each side denoted by $a \equiv |\mathcal{B}| = |\mathcal{S}| \in \{1, \ldots, \frac{n-1}{2}\}$, we use

$$\phi_\sigma(\boldsymbol{\omega}) \equiv p^B - (a+1)\,\sigma^B(\boldsymbol{\omega}) \text{ and } \gamma_\sigma(\boldsymbol{\omega}) \equiv p^S + (a+1)\,\sigma^S(\boldsymbol{\omega}), \tag{4}$$

where $\sigma^B(\boldsymbol{\omega})$ and $\sigma^S(\boldsymbol{\omega})$ are spacing estimators given by

$$\sigma^B(\boldsymbol{\omega}) \equiv \begin{cases} \frac{\theta_{(a+1)} - \theta_{(a+1+d(\boldsymbol{\omega}))}}{d(\boldsymbol{\omega})} & \text{if } d(\boldsymbol{\omega}) > 0, \\ \frac{1}{n+1} & \text{otherwise,} \end{cases} \quad \sigma^S(\boldsymbol{\omega}) \equiv \begin{cases} \frac{\theta_{[a+1+d(\boldsymbol{\omega})]} - \theta_{[a+1]}}{d(\boldsymbol{\omega})} & \text{if } d(\boldsymbol{\omega}) > 0, \\ \frac{1}{n+1} & \text{otherwise,} \end{cases} \tag{5}$$

and where $d(\boldsymbol{\omega})$ is the number of spacings to be included in the average, defined by

$$d(\boldsymbol{\omega}) \equiv \min\left\{\lceil n^c \rceil,\ \frac{n-1}{2} - a - 1\right\} \tag{6}$$

for some $c \in (0,1)$, where $\lceil n^c \rceil$ is the smallest integer greater than or equal to n^c. According to this definition, the hazard rate terms in the virtual types are estimated as the factor $a+1$ times the average of $d(\boldsymbol{\omega})$ spacings between adjacent types. For the virtual value, we use the average of the spacings between the types of the most recent $d(\boldsymbol{\omega})$ agents to exit on the buyer side, and for the virtual cost, we use the average of the spacings between the types of the most recent $d(\boldsymbol{\omega})$ agents to exit on the seller side. The number of spacings used depends on the parameter $c \in (0,1)$, which ensures that the number of spacings used increases with the number of agents, but at a slower rate. In that way, more information is used as more is available, but in the limit, the spacings used in the estimator are spacings between types that are "close" to the type of the marginal buyer or seller. We can then make use of the result of Loertscher and Marx (2020a) that ℓ times the average of a number of spacings close to the ℓ-th highest type converges in probability to $\frac{1-F(\theta_{(\ell)})}{f(\theta_{(\ell)})}$ and ℓ times the average of spacing close to the ℓ-th lowest type converges in probability to $\frac{F(\theta_{[\ell]})}{f(\theta_{[\ell]})}$.

Our prior-free asset market mechanism can therefore be summarized as $\langle \phi_\sigma, \gamma_\sigma, c \rangle$, defined by (4)–(6), where $c \in (0,1)$. Unpacking this, given a state $\boldsymbol{\omega}$ such that a active buyers and a active sellers remain, $\sigma^B(\boldsymbol{\omega})$ is the average spacing between the $d(\boldsymbol{\omega})$ highest types for nontrading buyer-side agents, i.e., the average spacing between the $d(\boldsymbol{\omega})$ highest types among

$$\hat{\theta} = \theta_{\left(\frac{n+1}{2}\right)}, \ldots, \theta_{(a+1)},$$

and analogously for the seller-side spacing $\sigma^S(\boldsymbol{\omega})$. The number of spacings to include is chosen so that σ^B is based only on spacings between types that are greater than the cutoff type and σ^S is based only on spacings between types that are less than the cutoff type.[14] In addition, the number of spacings to include is capped at n^c with $c \in (0,1)$ so that as n grows large, the number of spacings that can be included in the average grows large in absolute terms, but shrinks as a proportion of the total, i.e., as n goes to infinity, n^c goes to infinity, but n^c/n goes to zero. In addition, the share of spacings included in the average does not shrink too quickly in the sense that $\lim_{n\to\infty} \frac{\ln n}{n^c} = 0$. As a result, given our continuity assumptions on the distribution of types, the spacing estimators $\sigma^B(\boldsymbol{\omega})$ and $\sigma^S(\boldsymbol{\omega})$ are consistent estimators for the spacing between types close to $\theta_{(a)}$ and $\theta_{[a]}$, respectively, and we have the following:

Proposition 2. *The prior-free asset market mechanism $\langle \phi_\sigma, \gamma_\sigma, c \rangle$ with $c \in (0,1)$ is asymptotically optimal.*

[13] If $n = 3$, then the auction phase ends at $t = 0$ with 1 active buyer and 1 active seller, who trade at a price of $\hat{\theta}$, so in that case, the virtual type estimators do not play a role.

[14] During the auction phase at $t = 0$, the sets of active buyers and sellers have $\frac{n-1}{2}$ members, so $d(\boldsymbol{\omega_0}) = 0$. Thus, if $n \geq 5$, one has at most $\frac{n-3}{2}$ trades, even though having $\frac{n-1}{2}$ trades might increase the market maker's expected payoff.

Proof. See Appendix B.

While Proposition 2 shows that the prior-free asset market mechanism $\langle \phi_\sigma, \gamma_\sigma, c \rangle$ with $c \in (0,1)$ is asymptotically optimal, it leaves open the question as to what value c should take. The tradeoff is that with a smaller value of c, the hazard-rate estimator uses less information, but that information is more tightly clustered around the point where the hazard rate is being estimated; a larger value of c uses more information, but risks using information that is less relevant for the estimation of the hazard rate at the point under consideration. As discussed in Loertscher and Marx (2020a), using $c = 4/5$ gives the estimator that minimizes the mean square error of the estimator.[15]

While the prior-free asset market mechanism $\langle \phi_\sigma, \gamma_\sigma, c \rangle$ is optimal in the limit as the number of agents grows large, it can, of course, be used in settings with a finite, and even small, number of agents.

3.3. Partial clock auction implementation

It is possible to implement the direct mechanism derived above using a clock auction that applies to only one side of the market, thereby preserving the privacy of all trading agents on one side of the market. An ascending clock preserves the privacy of all traders who end up as buyers (i.e., whose types are larger than the market clearing price), and a descending clock auction preserves the privacy of all analogously defined sellers. In this way, the mechanism strikes a balance between eliciting the information required to operate and preserving the privacy of agents where possible. This is in a similar vein to the one-sided auction of Baranov et al. (2017), which induces losing bidders to reveal cost information about quantities that are no longer profitable, but generally does not require winning bidders to reveal their cost for the awarded quantities. As they remark, “the format strikes a balance between price discovery and privacy preservation” (Baranov et al., 2017, p. 5).

In contrast to a one-sided setup, in our asset market setup, agents who exit may still trade—an agent who exits an ascending (descending) clock auction may trade as a seller (buyer). Thus, the usual clock auction interpretation needs to be adapted to this environment. We refer to our auction implementation as a *partial clock auction.*

In an ascending partial clock auction, we first increase the buyer clock price from zero until only $\frac{n-1}{2}$ active buyers remain. At this point, the set of types that are less than or equal to $\hat{\theta}$, and their associated types, are observed. This allows us to initialize the state $\boldsymbol{\omega}_t = (\mathcal{B}_t, \mathcal{S}_t, p_t^B, p_t^S)$, with components as before, by setting $\mathcal{S}_0 \equiv \{i \in \mathcal{N} \mid \theta_i < \hat{\theta}\}$, $\mathcal{B}_0 \equiv \{i \in \mathcal{N} \mid \theta_i > \hat{\theta}\}$, and $p_0^B = p_0^S \equiv \hat{\theta}$. Then the clock auction defined previously can be used—as the buyer price increases, only the types of agents that exit from the buyer set are revealed. The privacy of all trading buyers is preserved. Analogously, one could start with a seller clock that descends from a price of one to identify the cutoff type and the set of agents with types greater than or equal to that type. Then the auction procedure preserves the privacy of all trading sellers.

While partial clock auctions endow agents with dominant strategies, they do not in general endow them with obviously dominant strategies. However, the truncated game that starts when the market clearing price is passed endows the remaining active agents with obviously dominant strategies from that point onwards because these remaining active agents are now essentially single-unit traders and the problem has become effectively one-sided.

The partial clock auction defined here is reminiscent of the incentive auction employed by the U.S. Federal Communications Commission in 2017 to purchase spectrum licenses from television broadcasters and then sell them to mobile wireless providers (see, e.g., Leyton-Brown et al., 2017; Milgrom, 2017). In that case, a reverse auction determined prices at which broadcasters were willing to relinquish their spectrum usage rights, and a forward auction determined the prices at which wireless carriers could purchase flexible-use wireless licenses, with an intermediate process for the reorganization and reassignment of spectrum licenses. Apart from the complexity inherent in the incentive auction due to interference constraints, conceptually, the key difference between the two formats is that here agents’ trading positions are determined endogenously, whereas the incentive auction required bidders to register as buyers or sellers.

The reason why a clock auction implementation of our asymptotically optimal prior-free mechanism is not possible is that the mechanism needs to learn the Walrasian price to be able to start the estimation and elimination procedure that we described algorithmically as consisting of two clock auctions. Consequently, if the Walrasian price is known before the mechanism starts, for example, because of historical data, then there is an asymptotically optimal mechanism that requires no prior information about distributions and that can be implemented using clock auctions. It dispenses with the first stage of our asymptotically optimal prior-free mechanism and directly starts with the estimation and elimination procedures using clock auctions for buyers and sellers that are initialized at the Walrasian price.

4. Performance in the small

Thus far, we have focused on the optimal mechanism in the large, and we have constructed a prior-free DIC-EIR mechanism that is optimal, asymptotically. Of course, any real-world mechanism would be applied in settings with a limited

[15] The choice of $c = 4/5$ minimizes the mean squared error of nearest neighbor estimators such as σ^B and σ^S (Silverman, 1986, Chapters 3 and 5.2.2). For related discussion, see Loertscher and Marx (2020a). While one can use all the spacings between exited agents, including those on the buyer and seller side (for a total of $n - 1 - 2a$ spacings) without interfering with incentive compatibility, this introduces a mechanical relation between ϕ and γ and so interferes with asymptotic optimality.

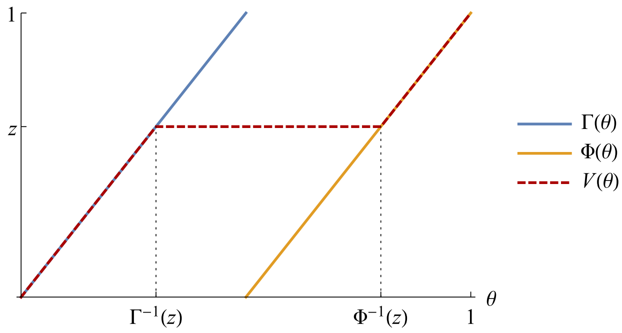

Fig. 2. Illustration of the ironed virtual type function with ironing parameter z.

number of traders, raising the question of how our prior-free mechanism performs with a finite, possibly small, number of agents. This is what we address in this section, beginning with the construction of the optimal mechanism in the small.

4.1. Optimal mechanism in the small

The *optimal* mechanism maximizes the market maker's expected profit, subject to DIC and EIR, given knowledge of the type distribution.[16] Because the underlying design problem is not regular even when the virtual type functions are monotone (see, e.g., Lu and Robert, 2001; Loertscher and Wasser, 2019), the optimal mechanism allocates on the basis of *ironed* virtual type functions (Myerson, 1981). Specifically, the ironed virtual type $V(\theta)$ of an agent of type θ is given as

$$V(\theta) \equiv \begin{cases} \Gamma(\theta) & \text{if } \Gamma(\theta) < z, \\ z & \text{if } \Phi(\theta) \le z \le \Gamma(\theta), \\ \Phi(\theta) & \text{if } z < \Phi(\theta), \end{cases} \tag{7}$$

where $z \in (0,1)$ is the unique number satisfying $F(\Gamma^{-1}(z)) = 1 - F(\Phi^{-1}(z))$.[17] The ironed virtual type function $V(\theta)$ is weakly increasing in θ, and it is equal to z for $\theta \in [\Gamma^{-1}(z), \Phi^{-1}(z)]$, which we refer to as the *ironing region*. We illustrate this in Fig. 2.

As Lu and Robert (2001) show, upon a vector of reports $\boldsymbol{\theta}$, the optimal mechanism induces the agents with the $q \in \{0, \dots, (n-1)/2\}$ lowest ironed virtual types to sell their units to the q agents with the highest ironed virtual types, with no trades occurring between two or more agents whose types are in the ironed region (put differently, whose ironed virtual types are equal to z).[18] That is, the quantity traded is

$$\min\left\{\frac{n-1}{2},\ \max\left\{\sum_{i\in\mathcal{N}} \mathbf{1}_{\Gamma(\theta_i)<z}, \sum_{i\in\mathcal{N}} \mathbf{1}_{\Phi(\theta_i)>z}\right\}\right\}.$$

The number z makes sure that the interim expected allocation of an agent who reports a type in the ironing region is 1. Because the outside option of each agent is to consume its endowment, the interim individual rationality constraints are satisfied with equality without further payments for these agents. Because these types are the worst off, the interim individual rationality constraint is satisfied for all types. Denoting by $\hat{s}_n(\theta)$ the interim expected final allocation of an agent who reports θ under this mechanism with n agents (see Appendix C for an analytic expression for $\hat{s}_n(\theta)$), the agent's interim expected quantity traded is $\hat{s}_n(\theta) - 1$, with $\hat{s}_n(\theta) < 1$ meaning that the agent is a seller in expectation and $\hat{s}_n(\theta) > 1$ meaning that the agent is a buyer in expectation. As usual, per-agent expected revenue can then be calculated as the expectation of the virtual type times the interim expected number of trades (see e.g., Börgers, 2015):

---

[16] Of course, DIC and EIR imply that the mechanism is also Bayesian incentive compatible and interim individually rational. Because the optimal mechanism involves randomization for a set of types $\boldsymbol{\theta}$ with positive measure, individual rationality in a direct mechanism can be evaluated *ex post* in two ways: after types have been reported and (i) before uncertainty due to randomization is resolved or (ii) after uncertainty due to randomization is resolved. That is, according to (i) agents make payments after types have been reported and obtain expected allocations that are conditional on the types reported, with the expectation being due to randomization. Interpreting EIR in the sense of (i), the optimal mechanism satisfying Bayesian incentive compatibility and interim individual rationality can also be implemented as a DIC-EIR mechanism.

[17] To see that such a number exists and is unique, note first that the left side increases in z and the right side decreases in z, so at most one such number exists. Observe then that both sides vary continuously in z and satisfy $F(\Gamma^{-1}(0)) = 0 < 1 - F(\Phi^{-1}(0))$ and $F(\Gamma^{-1}(1)) > 0 = 1 - F(\Phi^{-1}(1))$.

[18] For example, if $\ell < \frac{n-1}{2}$ of the agents have ironed virtual types less than z and $m < \frac{n-1}{2}$ of the agents have ironed virtual types greater than z, then the quantity traded is $\max\{\ell, m\}$. If, say, $\ell \ge m$, then all the agents with ironed virtual types less than z sell, all the agents with ironed virtual types greater than z buy, and $\ell - m$ randomly selected agents with ironed virtual types equal to z also buy.

(a) Revenue as a share of limit optimal revenue

n	Uniform	Beta(2,2)	Beta(4,1.5)
3	83.3%	85.2%	86.0%
5	94.4%	95.4%	95.7%
7	97.6%	98.1%	98.3%
9	98.8%	99.1%	99.2%

(b) Distributions

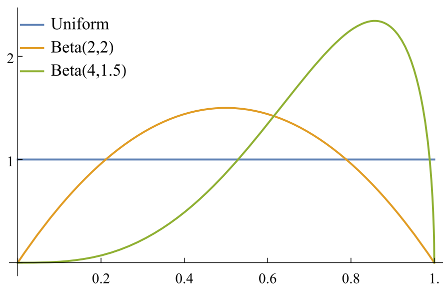

Fig. 3. Optimal per-agent expected revenue in the small as a fraction of the limit optimal revenue, $(p^B - p^S)Q^*$, for the uniform, Beta(2,2), and Beta(4,1.5) distributions, whose densities are displayed in panel (b).

$$\pi^O(n) \equiv \int_0^{\Gamma^{-1}(z)} \Gamma(x)(\hat{s}_n(x) - 1)dF(x) + \int_{\Phi^{-1}(z)}^{1} \Phi(x)(\hat{s}_n(x) - 1)dF(x).$$

As an illustration of performance in the small, in Fig. 3 we compare the optimal revenue with small numbers of agents to the limit optimal revenue for different distributions. As shown, profit with only 9 agents is approximately 99% of the limit optimal revenue for the distributions considered.

4.2. Prior-free mechanism in the small

With the benchmark of the optimal mechanism in hand, we are in a position to evaluate the performance of the prior-free mechanism in the small. The prior-free mechanism faces two challenges or obstacles: First, it must *estimate* the virtual types of the marginal agents, and related, to maintain incentive compatibility while achieving asymptotic optimality, never allow the median type to trade. Second, because the prior-free mechanism does not have an estimate of the type distribution over its entire domain, it cannot iron, implying that even absent estimation errors and trade prohibition for the median type, it cannot implement the optimal outcome.

To disentangle these effects and gain an understanding of their magnitudes, it is therefore useful to consider two intermediate mechanisms that are between our prior-free and the optimal mechanism: the *prior-free mechanism with no estimation error*, which has the same allocation rule as our prior-free mechanism if the prior-free mechanism made no estimation errors, and the *naive two-sided mechanism*, which ignores the asset market nature of the problem and "blindly" evaluates agents with types above the median as buyers and agents with types below the median as sellers and induces trade among buyers and sellers as long as the smallest virtual value of a buyer exceeds the largest virtual cost of a seller. Slightly more formally, the prior-free mechanism with no estimation error is defined by replacing the estimated ϕ and γ in the prior-free mechanism with the true Φ and Γ. This mechanism retains the structure of our prior-free mechanism, but eliminates the estimation error. The quantity traded of the naive two-sided mechanism is the $q \in \{0, \dots, (n-1)/2\}$ such that $\Phi(\theta_{(q)}) > \Gamma(\theta_{[q]})$. All agents i with $\theta_i \geq \theta_{(q)}$ obtain an additional unit and all agents i with $\theta_i \leq \theta_{[q]}$ sell their unit.

Fig. 4 illustrates these effects for the case with F uniform.[19] It shows the average number of trades and revenue per agent for the four mechanisms—our prior-free mechanism, the prior-free mechanism with no estimation error, the naive two-sided mechanism and the optimal mechanism—as the number of agents varies from 3 to 31. The effects of estimation errors are small as evidenced by the relatively small gap between the trades and the revenue under the prior-free mechanism and the prior-free mechanism with no estimation error. This contrasts with the substantial shortfall in revenue of the prior-free mechanisms relative to both the naive two-sided and the optimal mechanisms that arises from the prior-free mechanisms prohibiting the median agent from trading and from prohibiting all trades with $n = 3$.

Two things are noteworthy about the optimal mechanism with finite n, both of which relate to it making use of the ironed virtual type function. First, as shown in Fig. 4(b), for n small, the optimal mechanism achieves substantively higher revenue than the naive two-sided mechanism (for example, for $n = 3$, the naive two-sided mechanism achieves only about 62 percent of the revenue of the optimal mechanism), demonstrating the considerable revenue value of ironing. Second, as shown in Fig. 4(a), the probability that an agent trades in the optimal mechanism with n finite is larger than the probability that it trades in the continuum limit, i.e., it converges to $1/4$ from above, while for the three other mechanisms, the probability that an agent trades converges to $1/4$ from below. The reason is that the optimal mechanism uses agents with types in the ironed region to provide liquidity by serving as either buyers or sellers whenever the numbers of agents

[19] We calculate revenues for the naive two-sided mechanism based on threshold payments. If q units trade, then each trading seller is paid $\min\{\theta_{[q+1]}, \Gamma^{-1}(\Phi(\theta_{(q)}))\}$ and each trading buyer pays $\max\{\theta_{(q+1)}, \Phi^{-1}(\Gamma(\theta_{[q]}))\}$.

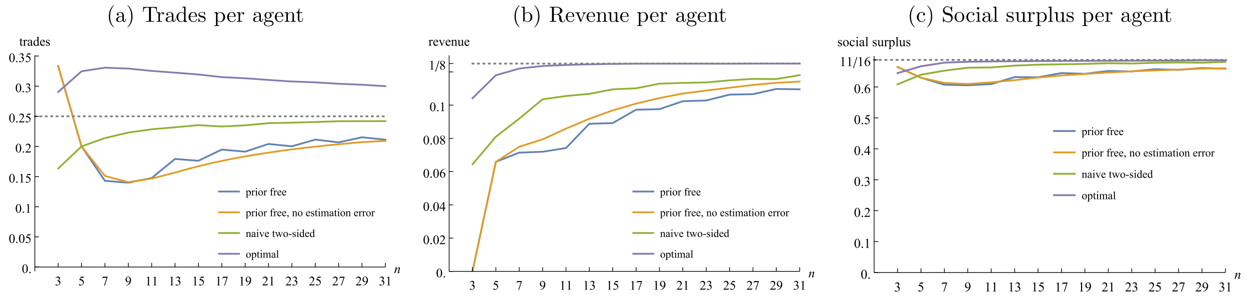

Fig. 4. Average number of trades, average revenue, and average social surplus per agent for 5000 simulated markets for varying numbers of agents as indicated. Prior-free and prior-free with no estimation error are identical for $n = 3$ and $n = 5$. Assumes uniformly distributed types on $[0, 1]$ and prior-free parameter $c = 4/5$. Social surplus is calculated as the sum of agents' types times their final quantity and has limiting value per agent of $\int_{p^S}^{p^B} x dF(x) + 2\int_{p^B}^{1} x dF(x)$, which is $11/16$ for uniformly distributed types.

with types below $\Gamma^{-1}(z)$ and above $\Phi^{-1}(z)$ are different.[20] This also means that the randomization that is involved in the optimal mechanism increases expected social surplus relative to the other mechanisms, in which no such randomization occurs, as shown in panel (c) of Fig. 4.

In fact, if f is symmetric in the sense that $f(\theta) = f(1-\theta)$ for all $\theta \in [0, 1]$, then one can show that, type profile by type profile, the optimal mechanism induces weakly more trade than the naive two-sided mechanism and strictly more than the prior-free mechanism with no estimation error. As shown in Appendix C, for f symmetric, $z = 1/2 = \Gamma(p^S) = \Phi(p^B)$. Because trade of $q \in \{0, \dots, \frac{n-1}{2}\}$ units in the naive two-sided mechanism requires that $\Phi(\theta_{(q)}) > \Gamma(\theta_{[q]})$, it must be that either $\theta_{[q]} < p^S$ or $\theta_{(q)} > p^B$, implying that the optimal mechanism trades at least q units and sometimes more.[21] Further, if more than $\frac{n-1}{2}$ agents have types below p^S or above p^B, then the optimal mechanism trades $\frac{n-1}{2}$ units (the maximum possible number of trades) even when $\Phi(\theta_{(\frac{n-1}{2})}) \le \Gamma(\theta_{[\frac{n-1}{2}]})$, in which case the naive two-sided mechanism has fewer than $\frac{n-1}{2}$ trades. Consequently, the quantity traded of the optimal mechanism is weakly larger than that of the naive two-sided mechanism. The quantity traded in the naive two-sided mechanism is one more than in the prior-free mechanism with no estimation error, which excludes one trade for which the virtual cost is less than the virtual value.[22] Thus, the optimal mechanism induces a larger quantity traded than the prior-free mechanism with no estimation error, type profile by type profile.

This discussion and analysis reveals three things. First, the asymptotically optimal prior-free mechanism tends to perform well even in the small insofar as its expected per-trader revenue converges quickly to that of the optimal mechanism. Second, deviations from the optimal mechanism relate to its sacrificing trades rather than estimation errors. As discussed, this could be remedied by augmenting the mechanism with a few natural and simple tweaks. Third, the advantage of the optimal mechanism even relative to the naive two-sided mechanism that makes use of knowledge of the true distribution but does not use ironing, provides motivation for incorporating ironing in constrained efficient, prior-free DIC-EIR mechanisms, which is a promising avenue for future research. As mentioned, this will require the mechanism to be able to extrapolate.[23]

5. Extensions

In this section, we provide extensions. In Section 5.1, we generalize the designer's objective to allow for positive weights on revenue and on social surplus. That is, we analyze a Ramsey pricing problem, which is of particular relevance in the context of the Laffer curve problem. In Section 5.2, we adapt the setup to allow agents to have decreasing marginal values. In Section 5.3, we extend the analysis beyond agents drawing their types from the same distribution to allow a priori heterogeneous groups of agents. In Section 5.4, we extend the setup to general endowments and demands.

[20] Because the probability that these numbers differ substantively remains large even for n relatively large, the probability that an agent trades remains much larger than 1/4 even for n in the order of 200 (simulations show that for $n = 201$, the probability of trade is 27.0%; for $n = 1,001$, it is 25.9%; and for $n = 10,001$, it is 25.3%). Despite the relatively slow convergence of the number of trades, the optimal expected revenue converges quickly to the limiting optimal expected revenue, as shown in Fig. 4(b).

[21] For example, if there are 5 agents with uniformly distributed types and $\boldsymbol{\theta} = \{0, .1, .3, .4, .6\}$, then there are two types less than $p^S = \frac{1}{4}$ and zero types greater than $p^B = \frac{3}{4}$. The naive two-sided mechanism trades one unit, but the optimal mechanism trades two units.

[22] The prior-free mechanism with no estimation error "starts" from the median trader and excludes the marginal seller and buyer from trading until and *including* the first pair where the seller's virtual cost is less than the buyer's virtual value.

[23] Random sampling mechanisms such as the one of Baliga and Vohra (2003) permit ironing but random sampling comes at the cost of giving up constrained efficiency.

5.1. Laffer curve and Ramsey pricing

As mentioned, the problem of the market maker who seeks to maximize its profit without knowing the agents' type distribution is equivalent to the problem of a tax authority that seeks to maximize its tax revenue without knowing the functional form of the Laffer curve (Laffer, 2004). To briefly illustrate this, letting $a \in \{0, \dots, \frac{n-1}{2}\}$ be the quantity traded, a mechanism that charges buyers the price $\theta_{(a+1)}$ and offers sellers the price $\theta_{[a+1]}$ nets a revenue of

$$R(a, \boldsymbol{\theta}) \equiv a\left(\theta_{(a+1)} - \theta_{[a+1]}\right).$$

Evidently, we have $R(0, \boldsymbol{\theta}) = 0 = R(\frac{n-1}{2}, \boldsymbol{\theta})$, which is reminiscent of the Laffer curve that has the same properties, typically expressed as a function of the tax rate rather than the quantity traded.[24] Just like for a market maker, a question of central interest for a tax authority is how much to distort trade away from the efficient level. Because the two problems are equivalent, our prior-free mechanism is also asymptotically optimal in the Laffer curve problem.

Of course, a social planner may not necessarily want to maximize tax revenue per se but rather, in the tradition of Ramsey (see e.g. Wilson, 1993), may be interested in maximizing a weighted average of tax revenue and social surplus, that is, maximizing the expected value of

$$\alpha(\text{revenue}) + (1-\alpha)(\text{social surplus}),$$

where $\alpha \in [0, 1]$ is a given parameter[25] and where the expectation is taken with respect to the distribution F.

In this formulation of the planner's objective, in the continuum limit, the planner's Ramsey problem is to choose Q to maximize

$$\alpha(F^{-1}(1-Q) - F^{-1}(Q))Q + (1-\alpha)\left[\int_{F^{-1}(1-Q)}^{1} 2x dF(x) + \int_{F^{-1}(Q)}^{F^{-1}(1-Q)} x dF(x)\right].$$

The maximizer Q^*_α is such that $1 - F(p^B_\alpha) = Q^*_\alpha = F(p^S_\alpha)$ and $\Phi_\alpha(p^B_\alpha) = \Gamma_\alpha(p^S_\alpha)$, where $\Phi_\alpha(\theta) \equiv (1-\alpha)\theta + \alpha\Phi(\theta)$ and $\Gamma_\alpha(\theta) \equiv (1-\alpha)\theta + \alpha\Gamma(\theta)$. The per-capita value of the objective in this limit is thus

$$\pi_\alpha \equiv \alpha(p^B_\alpha - p^S_\alpha)Q^*_\alpha + (1-\alpha)\left[\int_{p^B_\alpha}^{1} 2x dF(x) + \int_{p^S_\alpha}^{p^B_\alpha} x dF(x)\right].$$

Letting $a_n(\boldsymbol{\theta})$ be the quantity traded by a constrained efficient DIC-EIR mechanism when the type profile is $\boldsymbol{\theta}$ and the number of agents is n, the value of Ramsey objective ex post is

$$\alpha R(a_n(\boldsymbol{\theta}), \boldsymbol{\theta}) + (1-\alpha) SS(a_n(\boldsymbol{\theta}), \boldsymbol{\theta}),$$

where $SS(a_n(\boldsymbol{\theta}), \boldsymbol{\theta}) = \sum_{i=1}^{a_n(\boldsymbol{\theta})} 2\theta_{(i)} + \sum_{i=n-a_n(\boldsymbol{\theta})}^{a_n(\boldsymbol{\theta})+1} \theta_{(i)}$ is social surplus. A prior-free constrained efficient DIC-EIR mechanism is asymptotically optimal if

$$\lim_{n\to\infty} \frac{\mathbb{E}_{\boldsymbol{\theta}}[\alpha R(a_n(\boldsymbol{\theta}), \boldsymbol{\theta}) + (1-\alpha) SS(a_n(\boldsymbol{\theta}), \boldsymbol{\theta})]}{n} = \pi_\alpha.$$

Such a mechanism is obtained from our prior-free mechanism by simply augmenting the virtual type estimators with α as follows:

$$\phi(\boldsymbol{\omega}) = p^B - \alpha\,(a+1)\,\sigma^B(\boldsymbol{\omega}) \text{ and } \gamma(\boldsymbol{\omega}) = p^S + \alpha\,(a+1)\,\sigma^S(\boldsymbol{\omega}). \tag{8}$$

All results above then extend to this setting.

24 Of course, denoting by $\tau_a = \theta_{(a+1)} - \theta_{[a+1]}$ the tax rate that induces the quantity traded a, the Laffer curve is $a\tau_a$ and inherits the properties of $R(a, \boldsymbol{\theta})$.

25 For example, if all taxes are distortionary in the economy with every dollar of taxes raised reducing social surplus by $\lambda > 1$ units, then $\alpha = \lambda/(1+\lambda)$; see, for example, Norman (2004) or Loertscher et al. (2015).

5.2. Decreasing marginal values

Thus far, we assumed that an agent with type θ has a constant marginal value of θ for each of two units of the good. As we show here, it is straightforward to extend this setup to allow for decreasing marginal values.[26] To accomplish this, we adapt the approach of Loertscher and Marx (2022) for the bilateral trade setup, in which each agent is ex ante either a buyer or a seller, to our current setup in which that is not the case. Like there, we parameterize the extent to which agents' marginal values decrease for additional units with a commonly known parameter $\delta \in (0, 1]$ such that an agent with type θ has value θ for one unit of the good and total value for two units of the good of $\theta + \delta\theta$. Thus, the agent's marginal value for the first unit is θ and its marginal value for the second unit is weakly lower at $\delta\theta$.

Working within this extended setup is particularly tractable because the *virtual* value associated with agent i's second unit is simply $\delta\Phi(\theta_i)$. To see this, note that the distribution of $x = \delta\theta_i$ is $H(x) \equiv F(x/\delta)$ with density $h(x) \equiv \frac{1}{\delta} f(x/\delta)$ and support $[0, \delta]$, and so agent i's virtual value for its second unit is

$$\delta\theta_i - \frac{1 - H(\delta\theta_i)}{h(\delta\theta_i)} = \delta\theta_i - \delta\frac{1 - F(\theta_i)}{f(\theta_i)} = \delta\Phi(\theta_i).$$

With a slight abuse of notation, the optimal outcome in the limit case can then be defined in terms of prices p^B_δ and p^S_δ satisfying market clearing,

$$1 - F(p^B_\delta/\delta) = Q^*_\delta = F(p^S_\delta), \tag{9}$$

and equating marginal revenue and marginal cost,[27]

$$\delta\Phi(p^B_\delta/\delta) = \Gamma(p^S_\delta). \tag{10}$$

For example, for F uniform, we have

$$p^B_\delta = \frac{\delta(2+\delta)}{2(1+\delta)} \text{ and } p^S_\delta = \frac{\delta}{2(1+\delta)}.$$

To give a particular example, if $\delta = \frac{1}{2}$, then we have $p^B_\delta = \frac{5}{12}$ and $p^S_\delta = \frac{1}{6}$, as compared with the corresponding prices of $\frac{3}{4}$ and $\frac{1}{4}$ when $\delta = 1$. As this illustrates, a diminished marginal value for the second unit requires a reduction in the price to buyers (because agents' willingnesses to pay as a buyer are reduced) and a reduction in the price to sellers to balance the reduced demand by buyers.

As in the model with constant marginal values, it remains the case that as n goes to infinity, it is a probability zero event that there are trades by agents with types in the ironing region, and so the methodology described above for estimating the virtual values and virtual costs continues to deliver an asymptotically optimal prior-free mechanism.

5.3. Price discrimination among heterogeneous groups

We now explore the extent to which the analysis and mechanism extend to allow for ex ante heterogeneous groups of traders, at the cost, of course, of sacrificing constrained efficiency across groups because of the discrimination.[28] We do this for the model with constant marginal values.

In one-sided allocation problems such as auctions or procurement auctions and in two-sided settings with agents that always trade either as buyers or sellers, the worst-off types are known a priori (the lowest possible type for a buyer and the highest possible type for a seller). This permits prior-free, asymptotically optimal mechanisms because the designer knows a priori where to start the elimination procedure required for estimation.[29] In contrast, as seen above, in an asset market, individual rationality constraints typically bind in the interior of the type space. With a homogeneous population, our asymptotically optimal prior-free asset market mechanism initiates the elimination procedure at the median, which is prior-free. We now show that this procedure extends to environments with ex ante heterogeneous groups, provided the heterogeneity among groups with respect to their relative sizes and distributions is not too accentuated.

Formally, for the purposes of this section, we assume that there are m groups of agents. Let $\mathcal{M} \equiv \{1, \ldots, m\}$, and for each $j \in \mathcal{M}$, let n_j denote the number of agents in group j. We assume that each agent in group j draws its type independently from the distribution F_j, where, as before, F_j is continuously differentiable with positive density f_j on $[0, 1]$. Let Φ_j and Γ_j

[26] On the asymptotic efficiency of double auctions in a setting with decreasing marginal values and increasing marginal costs, see, for example, Cripps and Swinkels (2006) or Loertscher and Mezzetti (2021).

[27] To see this, note that the market maker's problem can be expressed as $\max_Q \left(H^{-1}(1-Q) - F^{-1}(Q)\right) Q$, which, using $H^{-1}(1-Q) = \delta F^{-1}(1-Q)$, can be rewritten as $\max_Q \left(\delta F^{-1}(1-Q) - F^{-1}(Q)\right) Q$. The expression in (10) then follows from the first-order condition and (9).

[28] As will become clear, within each group the mechanism remains constrained efficient in the sense that only the highest value agents buy and only the lowest value agents sell.

[29] This is also why for problems like these the asymptotically optimal prior-free mechanisms can be implemented via clock auctions (Loertscher and Marx, 2020a).

denote the corresponding virtual value and virtual cost functions, assumed to be increasing on [0, 1]. When analyzing the limits as the numbers of agents goes to infinity, we assume that those numbers increase in fixed proportion to one another.

Let $(\overline{p}_j^B)_{j\in\mathcal{M}}$ and $(\overline{p}_j^S)_{j\in\mathcal{M}}$ denote the optimal prices in the limit economy when all groups or markets are integrated, which satisfy, for all $j \in \mathcal{M}$,

$$\Phi_j(\overline{p}_j^B) = \Gamma_j(\overline{p}_j^S) \equiv V \tag{11}$$

and the market clearing condition

$$\sum_{j\in\mathcal{M}} \mu_j(1 - F_j(\overline{p}_j^B)) = \sum_{j\in\mathcal{M}} \mu_j F_j(\overline{p}_j^S), \tag{12}$$

where $\mu_j \equiv \frac{n_j}{\sum_{i\in\mathcal{M}} n_i} > 0$ is the probability that a randomly selected trader draws its type from the distribution F_j.

5.3.1. Prior-free mechanism with heterogeneous groups

We now define a prior-free asset market mechanism for the setup with groups. As in the case without heterogeneous groups, the mechanism is direct, elicits reports from all agents, and identifies the cutoff type $\hat{\theta}$ as the median report. We initialize buyer and seller prices for each group $j \in \mathcal{M}$ as $p_0^{S_j} = p_0^{B_j} = \hat{\theta}$.[30] We denote the estimated virtual value for the marginal buyer in group j by ϕ_j and the estimated virtual cost for the marginal seller in group j by γ_j, where the estimation procedure is as above, except that the estimates for a given agent are based only on the reports of other agents from the same group. For each $t \in \{0, 1, \ldots\}$, follow an iterative process similar to that defined above for the setup without groups, but with a state amended to include the sets of active buyers and sellers in each group and the buyer and seller prices in each group. In step 1 of the auction phase, if the minimum value of $\phi_j(\boldsymbol{\omega}_t)$ across all groups j with at least one remaining active buyer is greater than or equal to the maximum value of $\gamma_j(\boldsymbol{\omega}_t)$ across all groups j with at least one remaining active seller, then the auction phase ends. If not, then the buyer price for the group with the lowest virtual value is increased to be the minimum value among the remaining active buyers in that group, and that buyer is removed from the set of active buyers in that group, and the seller price for the group with the greatest virtual cost is decreased to be the maximum cost among remaining active sellers in that group, and that seller is removed from the set of active sellers in that group. Then t is incremented by 1 and the procedure continues with an updated state.

5.3.2. Asymptotical optimality with heterogeneous groups

Denote by p^W the market clearing price in the limit large economy with heterogeneous groups. That is, p^W is such that $\sum_{j\in\mathcal{M}} \mu_j(1 - F_j(p^W)) = \sum_{j\in\mathcal{M}} \mu_j F_j(p^W)$.

Proposition 3. *If for all $j \in \mathcal{M}$,*

$$\Phi_j(p^W) < V < \Gamma_j(p^W), \tag{13}$$

then the prior-free asset market mechanism for groups defined above is asymptotically optimal.

Condition (13) implies that agents with types equal or close to the market clearing price in the large economy do not trade in the optimal mechanism independently of the group to which they belong. It prevents "excessive" heterogeneity across groups in the sense that, even though there is discrimination across groups, the discrimination effect is weak enough that an agent with type p^W neither trades in the agent's own group when that group is treated as a standalone market nor in the integrated market. One can therefore use the price p^W to initiate estimation and elimination without eliminating the "wrong" agents. As for the case of homogeneous agents, one can apply Loertscher and Marx (2020a, Theorem 1) to show that this mechanism is asymptotically optimal, which is why we do not include an independent proof of Proposition 3.

The optimality condition in the large-market limit in standalone market j is that the prices p_j^B and p_j^S satisfy both

$$\Phi_j(p_j^B) = \Gamma_j(p_j^S) \equiv V_j \tag{14}$$

and the market-clearing condition

$$\sum_{j\in\mathcal{M}} \mu_j(1 - F_j(p_j^B)) = \sum_{j\in\mathcal{M}} \mu_j F_j(p_j^S). \tag{15}$$

Conditions (14) and (15) holding for every standalone market j, we let

$$\underline{V} = \min_{j\in\mathcal{M}}\{V_j\} \text{ and } \overline{V} = \max_{j\in\mathcal{M}}\{V_j\}.$$

[30] As before, the median report is well defined if $n = \sum_j n_j$ is odd, but one can easily accommodate an even number of agents n with $n \geq 4$ by defining $\hat{\theta}_1 \equiv \theta_{(n/2+1)}$ and $\hat{\theta}_2 \equiv \theta_{(n/2)}$, where $\hat{\theta}_1 \leq \hat{\theta}_2$, and initializing the auction phase with $p_0^{S_j} \equiv \hat{\theta}_1$ and $p_0^{B_j} \equiv \hat{\theta}_2$.

Lemma 1. $V \in [\underline{V}, \overline{V}]$.

Proof. See Appendix B.

Lemma 1 implies that a sufficient condition for (13) to hold is what we refer to as *Virtual Walrasian gap overlap*, which is stated formally as Assumption 1[31]:

Assumption 1. For all $j \in \mathcal{M}$, $\Phi_j(p^W) < \underline{V}$ *and* $\overline{V} < \Gamma_j(p^W)$.

In the model with ex ante homogeneous agents or with moderate heterogeneity across groups of agents, insisting that a mechanism satisfy DIC for any number of agents does not come at a cost in terms of the revenue that it generates asymptotically. In contrast, if heterogeneity is large, a tradeoff can arise between DIC and asymptotic optimality. In that case, relaxing the notion of incentive compatibility by, for example, replacing it by approximate strategy proofness in the large may have the benefit of permitting asymptotic optimality. While this seems a worthwhile avenue for future research, its pursuit is beyond the scope of the present paper.

5.4. General endowments and demands

We now show how the setup and analysis extends beyond the setting in which each agent has an endowment of one unit and a maximum demand for two.

5.4.1. Setup and optimal outcome for the general case

In the generalized setup, each agent $i \in \mathcal{N}$ has endowment r_i and a maximum demand k_i, where $0 \leq r_i < k_i$ for all $i \in \mathcal{N}$, with $r_i > 0$ for at least one $i \in \mathcal{N}$. To define the optimal outcome requires some additional notation. Define $\boldsymbol{\theta}^B$ to be the vector that replicates each agent i's type $k_i - r_i$ times:

$$\boldsymbol{\theta}^B \equiv (\overbrace{\theta_1, \ldots, \theta_1}^{k_1 - r_1 \text{ times}}, \ldots, \overbrace{\theta_n, \ldots, \theta_n}^{k_n - r_n \text{ times}}),$$

and define $\boldsymbol{\theta}^S$ to be the vector that replicates each agent i's type r_i times:

$$\boldsymbol{\theta}^S \equiv (\overbrace{\theta_1, \ldots, \theta_1}^{r_1 \text{ times}}, \ldots, \overbrace{\theta_n, \ldots, \theta_n}^{r_n \text{ times}}).$$

In the optimal mechanism in the large, q units are traded, where q is the largest index in

$$\left\{0, 1, \ldots, \min\left\{\sum_{i \in \mathcal{N}} (k_i - r_i), \ \sum_{i \in \mathcal{N}} r_i\right\}\right\}$$

such that

$$\Phi(\theta^B_{(q)}) \geq \Gamma(\theta^S_{[q]}).$$

Each agent i with type greater than $\theta^B_{(q)}$ purchases $k_i - r_i$ units, and so satisfies its maximum demand, and each agent i with type less than $\theta^S_{[q]}$ sells r_i and so consumes zero. The agent with type $\theta^B_{(q)}$ purchases $q - \sum_{i\,:\,\theta_i > \theta^B_{(q)}} (k_i - r_i)$ units and the agent with type $\theta^S_{[q]}$ sells $q - \sum_{i\,:\,\theta_i < \theta^S_{[q]}} r_i$ units.

To specify the threshold payments, first, we say that the seller side is the long side of the market if

$$\sum_{i\,:\,\theta_i \geq \theta^B_{(q)}} (k_i - r_i) \leq \sum_{i\,:\,\theta_i \leq \theta^S_{[q]}} r_i,$$

and otherwise we say that the buyer side is the long side of the market. If the seller side is the long side of the market, then all buyers pay $\theta^B_{(q+1)}$, which is less than $\theta^B_{(q)}$ when the seller side is the long side of the market. If the buyer side is the long side of the market, then all sellers are paid $\theta^S_{[q+1]}$, which is greater than $\theta^S_{[q]}$ when the buyer side is the long side. It then remains to describe the threshold payments for agents on the long side of the market, which are Vickrey prices (see, e.g., Krishna, 2002, Chapter 12). If the buyer side is the long side, then buyers' payments are the payments that would arise in a Vickrey auction to sell q units to the agents with types greater than or equal to $\theta^B_{(q)}$ with a reserve (minimum price)

[31] Assumption 1 and the condition that (13) holds for all $j \in \mathcal{M}$ can be relaxed by stipulating that they hold for some a priori *known* quantile q_j in every market j.

equal to $\max\{\theta_i \mid \theta_i < \theta^B_{(q)}\}$. If the seller side is the long side, then payments to the sellers are the payments that would arise in a Vickrey auction to purchase q units from the agents with types less than or equal to $\theta^S_{[q]}$ with a reserve (maximum price) equal to $\min\{\theta_i \mid \theta_i > \theta^S_{[q]}\}$. For a review of the details of Vickrey pricing, see Appendix A or Loertscher and Marx (2020b).

5.4.2. Asymptotic optimality

In the generalized setup, asymptotics and asymptotic optimality are then captured, as in Gresik and Satterthwaite (1989), by looking at η-fold replicas of the economy characterized by $(k_i, r_i)_{i\in\mathcal{N}}$ and letting η go to infinity. More specifically, fixing the initial number of agents n, each replica adds an additional set of n agents characterized by $(k_i, r_i)_{i\in\mathcal{N}}$, with each agent independently drawing its type from F. Given $\eta \in \{1, 2, \dots\}$, an η-fold replica of this initial economy has ηn agents, where η have characteristics (k_1, r_1), η have (k_2, r_2), etc., up to η that have (k_n, r_n), each with an independently drawn type. We let $\mathcal{N}_\eta$ denote the set of agents in the η-th replica. In this setup, for all replicas, the market maker's beliefs about the type space amount to beliefs over F, which is the same as for a single replica.

Let

$$\bar{r} = \sum_{i\in\mathcal{N}} \frac{r_i}{n} \text{ and } \bar{d} = \sum_{i\in\mathcal{N}} \frac{k_i - r_i}{n}$$

be per-agent supply and net demand, respectively, in a given replica. Then, in the limit as η goes to infinity, the profit-maximizing mechanism is characterized by prices p^B and p^S such that agents with $\theta \in (p^S, p^B)$ do not trade, agents with $\theta_i \geq p^B$ buy $k_i - r_i$, and agents with $\theta_j \leq p^S$ sell r_j, with (p^S, p^B) satisfying

$$\Phi(p^B) = \Gamma(p^S) \quad \text{and} \quad \bar{d}(1 - F(p^B)) = \bar{r}F(p^S).$$

It follows that in the limit, the market maker's per-agent expected profit is

$$\lim_{\eta\to\infty} \frac{1}{\eta n} \left(\sum_{i\in\mathcal{N}_\eta} \mathbb{E}\left[p^B(k_i - r_i) \mid \theta_i \geq p^B\right] - \sum_{i\in\mathcal{N}_\eta} \mathbb{E}\left[p^S r_i \mid \theta_i \leq p^S\right] \right).$$

Now that we have identified the optimal outcome and characterized its asymptotic properties, we can define what it means for a mechanism to be asymptotically optimal. We say that an incentive compatible, individually rational mechanism defined for an η-fold replica is *asymptotically optimal* if the ratio of the market maker's expected profit in the mechanism relative to its expected profit in the optimal mechanism converges to 1 in probability as the number of replicas goes to infinity.

5.4.3. Incentive compatible asset market mechanism

The iterative steps in the auction phase of the asset market mechanism generalize as follows:

- Step 1: If $\phi(\boldsymbol{\omega}_t) \geq \gamma(\boldsymbol{\omega}_t)$ or $|\mathcal{B}| \leq 1$ or $|\mathcal{S}| \leq 1$, then the estimation phase ends. Otherwise, proceed to step 2.
- Step 2: Let

$$\boldsymbol{\omega}_{t+1} = \begin{cases} (\mathcal{B}_t \backslash \{j^B\},\ \mathcal{S}_t,\ \theta_{j^B},\ p^S_t) & \text{if } \sum_{i\in\mathcal{B}_t}(k_i - r_i) \geq \sum_{i\in\mathcal{S}_t} r_i, \\ (\mathcal{B}_t,\ \mathcal{S}_t \backslash \{j^S\},\ p^B_t,\ \theta_{j^S}) & \text{otherwise,} \end{cases}$$

  where j^B and j^S are defined by $\theta_{j^B} = \min_{i\in\mathcal{B}_t}\{\theta_i\}$ and $\theta_{j^S} = \max_{i\in\mathcal{S}_t}\{\theta_i\}$, and proceed to step 3.
- Step 3: Increment t by 1 and return to step 1.

When the estimation phase ends, with final state $\boldsymbol{\omega}^B = (\mathcal{B}, \mathcal{S}, p^B, p^S)$, then Q units trade, where

$$Q = \min\left\{ \sum_{i\in\mathcal{B}} (k_i - r_i),\ \sum_{i\in\mathcal{S}} r_i \right\}.$$

If $\sum_{i\in\mathcal{B}}(k_i - r_i) = \sum_{i\in\mathcal{S}} r_i$, then each buyer $i \in \mathcal{B}$ receives $k_i - r_i$ units and pays $p^B(k_i - r_i)$ and each seller $i \in S$ provides r_i units and is paid $p^S r_i$. If $\sum_{i\in\mathcal{B}}(k_i - r_i) < \sum_{i\in\mathcal{S}} r_i$, then each buyer $i \in \mathcal{B}$ receives $k_i - r_i$ units and pays $p^B(k_i - r_i)$ and the market maker procures Q units from sellers in $\mathcal{S}$ using a Vickrey auction with reserve p^S. If $\sum_{i\in\mathcal{B}}(k_i - r_i) > \sum_{i\in\mathcal{S}} r_i$, then each seller $i \in \mathcal{S}$ provides r_i units and is paid $p^S r_i$ and the market maker sells the Q units purchased from the sellers to the buyers in $\mathcal{B}$ using a Vickrey auction with reserve p^B.

With this generalization, Proposition 1 continues to hold. The proof follows along the lines of the proof of Proposition 1, noting that the nature of clock and Ausubel auctions implies that in the auction phase every agent has a dominant strategy to bid sincerely. Asymptotic optimality holds using the same virtual type estimators as defined above.

5.4.4. Heterogeneous groups

We know briefly discuss how the model with heterogeneous k_i and r_i can be extended to allow for groups of agents that are heterogeneous with respect to their distributions as in Section 5.3. Just like there, we let $\mathcal{M}$ denote the set of groups with every group $j \in \mathcal{M}$ being characterized by a distribution F_j that exhibits increasing virtual type functions. For agent $i \in \mathcal{N}$ in the initial economy, we denote by $j(i)$ the group it belongs to. Thus, the initial economy is described by $(k_i, r_i, j(i))_{i \in \mathcal{N}}$, and we study η-fold replicas of this economy, in which the set of groups $\mathcal{M}$ is fixed. That is, in the η-fold replica, we have ηn agents, η of which are described ex ante as $(k_1, r_1, j(1))$, $(k_2, r_2, j(2))$, and so on.

Letting $\mathcal{M}_j = \{i \in \mathcal{N} \mid j(i) = j\}$ be the set of agents in group $j \in \mathcal{M}$ in the initial economy, we can express the net demand and the supply of agents in the initial economy in group j as

$$\overline{d}_j = \sum_{i \in \mathcal{M}_j} (k_i - r_i) \quad \text{and} \quad \overline{r}_j = \sum_{i \in \mathcal{M}_j} r_i.$$

Consequently, in the limit economy, that is, as $\eta \to \infty$, the market-clearing condition becomes

$$\sum_{j \in \mathcal{M}} \mu_j \overline{d}_j (1 - F_j(\overline{p}_j^B)) = \sum_{j \in \mathcal{M}} \mu_j \overline{r}_j F_j(\overline{p}_j^S). \tag{16}$$

With these adjustments, our prior results continue to hold. In particular, if for all groups j, $\Phi_j(p^W) < V < \Gamma_j(p^W)$ holds, where p^W is the market clearing price in the limit economy and V is the virtual type of the marginal traders in the optimal mechanism in the limit, then the appropriately adjusted prior-free mechanism with groups is asymptotically optimal.

6. Conclusion

The advent of the gig economy—Uber, AirBnB, Nerdify, etc.—has rendered many markets that used to be two-sided into asset markets, with agents' trading positions being determined endogenously. Because these asset markets are typically organized by profit-seeking intermediaries who operate in new environments with demand and supply functions that are not known, they give renewed salience to the problem studied here: developing prior-free asset market mechanisms that are asymptotically optimal.

Under the assumption that the underlying distributions are regular, we derive a prior-free, asymptotically optimal mechanism that allocates the quantity traded efficiently when all agents draw their types from the same distribution. With heterogeneous type distributions that differ across commonly known groups of agents, the mechanism generalizes, provided that the heterogeneity is not excessive. A major challenge in asset markets is identifying the types who do not trade in the profit-maximizing mechanism; these types are typically in the interior of the type space, making the problem different from one-sided or two-sided allocations problems (that is, auctions or double auctions), where it is a priori clear which buyer and seller types do not trade. Although privacy preservation via clock auctions is not possible in asset markets, our asymptotically optimal asset market mechanism can algorithmically be described as a clock auction and permits partial clock implementation. Further research on asset market problems with heterogeneous goods seems particularly relevant and promising.

Declaration of competing interest

None.

Appendix A. Vickrey auction

This appendix draws heavily from Loertscher and Marx (2020b, Appendix A.1).

We first review a Vickrey auction that is used to *sell* q units with a reserve of p to a set of agents $\mathcal{A}$ with types that are all greater than or equal to p. The Vickrey auction allocates to each agent $i \in \mathcal{A}$ a quantity of $\overline{q}_i$ units, where

$$\overline{q}_i \equiv \min\left\{k_i - r_i,\ \max\{0,\ q - \sum_{\ell \in \mathcal{A}\,:\,\theta_\ell > \theta_i} (k_\ell - r_\ell)\}\right\}.$$

Thus, agent i is allocated units whenever there remain units available from the total of q after the net demands of agents in $\mathcal{A}$ with types greater than θ_i have been satisfied, up to a maximum of $k_i - r_i$ units.

The amount paid by agent i is based on an individualized price vector consisting of p followed by the types of the other agents in $\mathcal{A}$, in increasing order by their types. Specifically, letting $\overline{\theta}_{[\ell]}^{-i}$ denote the type associated with the ℓ-th lowest type element of $\{\theta_j\}_{j \in \mathcal{A} \setminus \{i\}}$, we define:

$$\overline{\mathbf{p}}^i \equiv (p, \overline{\theta}_{[1]}^{-i}, \ldots, \overline{\theta}_{[|\mathcal{A}|-1]}^{-i}).$$

	p	$<$	θ_3	$<$	θ_2	$<$	θ_1
$k_i - r_i$			1		2		1

Fig. 5. Example of net demand by agents with types greater than p.

For example, in the setup of Fig. 5, if $\mathcal{A} = \{1, 2, 3\}$ and $q = 3$, then the 3 units are allocated first to the higher-type agents so that $\overline{q}_1 = 1$, $\overline{q}_2 = 2$, and $\overline{q}_3 = 0$. Further, the individualized price vectors are simply the vectors consisting of p followed by the types of the other agents in $\mathcal{A}$, in increasing order: $\overline{\mathbf{p}}^1 = (p, \theta_3, \theta_2)$, $\overline{\mathbf{p}}^2 = (p, \theta_3, \theta_1)$, and $\overline{\mathbf{p}}^3 = (p, \theta_2, \theta_1)$.

Agent i's total payment is determined by applying the prices in $\overline{\mathbf{p}}^i$ to traunches of agent i's units. Agent i pays p for the first $\overline{b}^i_0$ units that it purchases, pays $\overline{\theta}^{-i}_{[1]}$ for the next $\overline{b}^i_1$ units, etc. The traunches of units are defined so that each agent pays an amount equal to the externality that it exerts on the other agents. That is, letting $d^{-i}_{[\ell]}$ be the net demand (capacity minus endowment) of the agent associated with type $\overline{\theta}^{-i}_{[\ell]}$, we define $\overline{b}^i_0$ to be the maximum quantity (up to $\overline{q}_i$) out of q that can be allocated to agent i without affecting the allocation of the other agents in $\mathcal{A}$,

$$\overline{b}^i_0 \equiv \max\left\{0,\ \min\left\{\overline{q}_i,\ q - \sum_{\ell=1}^{|\mathcal{A}|-1} d^{-i}_{[\ell]}\right\}\right\}, \tag{17}$$

and for $\ell \in \{1, \dots, |\mathcal{A}| - 1\}$, we define $\overline{b}^i_\ell$ iteratively to be the additional quantity (up to $\overline{q}_i - \sum_{t=0}^{\ell-1} \overline{b}^i_t$) that can be allocated to agent i by imposing an externality on the ℓ-th lowest type agent in $\mathcal{A}$,

$$\overline{b}^i_\ell \equiv \max\left\{0,\ \min\left\{\overline{q}_i - \sum_{t=0}^{\ell-1} \overline{b}^i_t,\ d^{-i}_{[\ell]}\right\}\right\}. \tag{18}$$

Applying the prices in $\overline{\mathbf{p}}^i$ to the traunches for agent i, we get a total payment for agent i of $\overline{\mathbf{p}}^i \cdot \overline{\mathbf{b}}^i$.

In the example of Fig. 5, assuming $q = 3$, we have $\mathbf{b}^1 = (0, 1, 0)$ and $\mathbf{b}^2 = (1, 1, 0)$. Thus, agent 1 pays θ_3 for its unit and agent 2 pays p for its first unit and θ_3 for its second unit. Because there are 3 units for sale and only two units demanded by agents other than agent 2, agent 2 is guaranteed to trade one unit and so pays the reserve for that unit. Agent 2's consumption of that unit has no impact on the trades of the others. However, agent 2's purchase of a second unit precludes agent 3 from trading, so agent 2 pays θ_3 for its second unit. Similarly, agent 1's purchase of a unit precludes agent 3 from trading, so agent 1 pays θ_3 for its unit.

Now turn to the Vickrey auction to purchase q units with reserve p from a set of agents $\mathcal{A}$ with types less than or equal to p. In this case, the quantity purchased from agent i is

$$\underline{q}_i \equiv \min\left\{r_i,\ \max\{0,\ q - \sum_{\ell \in \mathcal{A}\,:\,\theta_\ell < \theta_i} r_\ell\}\right\},$$

so that agent i sells units whenever there remain units demanded after the agents ranked below agent i have sold their endowments, up to a maximum of r_i units. Analogously to the buyer side case, let $\underline{\theta}^{-i}_{(\ell)}$ be the type associated with the ℓ-highest element of $\{\theta_j\}_{j \in \mathcal{A}\setminus\{i\}}$.

The individualized price vector for agent i is

$$\underline{\mathbf{p}}^i \equiv (p, \underline{\theta}^{-i}_{(1)}, \dots, \underline{\theta}^{-i}_{(|\mathcal{A}|-1)}).$$

The payment made to agent i is determined by applying the prices in $\underline{\mathbf{p}}^i$ to traunches of units $\underline{\mathbf{b}}^i$ defined as in (17)–(18), but with $\overline{q}_i$ replaced by $\underline{q}_i$ and $d^{-1}_{[\ell]}$ replaced by the supply of the agent with type $\underline{\theta}^{-i}_{(\ell)}$. Then each agent $i \in \mathcal{A}$ is paid $\underline{\mathbf{p}}^i \cdot \underline{\mathbf{b}}^i$.

Appendix B. Proofs

The following corollary extends Lu and Robert (2001, Theorem 3) to the limit case:

Corollary B.1. *In the limit as n goes to infinity, with probability one, the optimal outcome requires each agent with type less than p^S to sell one unit and each agent with type greater than p^B buy one unit.*

Proof of Corollary B.1. Let $x^* = p^S$ and note that because $y(x^*) = \Phi^{-1}(\Gamma(x^*))$, we have $y(x^*) = p^B$, giving us an ironing range of $[p^S, p^B]$. Define the allocation rule $\mathbf{s}^*$ so that when the empirical distribution of types is equal to F, which occurs with probability one in the limit, then

$$s_i^*(\boldsymbol{\theta}) = \begin{cases} 0 & \text{if } \theta_i < p^S, \\ 2 & \text{if } \theta_i > p^B, \\ 1 & \text{otherwise,} \end{cases}$$

which is feasible by (3). For other type realizations, let the allocation rule be (arbitrarily) such that no agent trades. Clearly, under this allocation rule and ironing region, the interim expected trade for agents with types in the ironing region is zero. With probability one, the empirical distribution is equal to F, so the optimality of the allocation rule given the ironing region follows because for all $\theta_i > p^B$ and $\theta_j < p^S$, $\Phi(\theta_i) > \Gamma(\theta_j)$, and so $\mathbf{s}^*$ solves

$$\max_{\mathbf{s}} \mathbb{E}_{\boldsymbol{\theta}} \left[\lim_{n\to\infty} \sum_{i=1}^{n} \left(\Phi(\theta_i) \mathbf{1}_{s_i(\boldsymbol{\theta})=2} - \Gamma(\theta_i) \mathbf{1}_{s_i(\boldsymbol{\theta})=0} \right) \right],$$

subject to market clearing. ■

Proof of Proposition 2. Given Proposition 1, we can assume truthful reporting in the prior-free asset market mechanism.

Loertscher and Marx (2020a, Theorem 1) states that if F is continuously differentiable with positive density f and compact support and $c \in (0,1)$, as holds in our setup, then as $n \to \infty$,

$$\Pr\left[\sup_{\ell \in \{1,\ldots,n-n^c\}} \left| \ell \frac{\theta_{(\ell)} - \theta_{(\ell+n^c)}}{n^c} - \frac{1 - F(\theta_{(\ell)})}{f(\theta_{(\ell)})} \right| > \varepsilon \right] \to 0 \tag{19}$$

and

$$\Pr\left[\sup_{\ell \in \{1,\ldots,n-n^c\}} \left| \ell \frac{\theta_{[\ell+n^c]} - \theta_{[\ell]}}{n^c} - \frac{F(\theta_{[\ell]})}{f(\theta_{[\ell]})} \right| > \varepsilon \right] \to 0. \tag{20}$$

At iteration $t \in \{0, 1, \ldots\}$ of the prior-free asset market mechanism, we have state $\boldsymbol{\omega}_t$ with $a_t \equiv |\mathcal{B}_t| = |\mathcal{S}_t| = \frac{n-1}{2} - t$. Further, $\phi(\boldsymbol{\omega}_0) = \hat{\theta} - \left(\frac{n-1}{2} + 1\right) \frac{1}{n+1} = \hat{\theta} - \frac{1}{2}$ and $\gamma(\boldsymbol{\omega}_0) = \hat{\theta} + \frac{1}{2}$, so $\phi(\boldsymbol{\omega}_0) < \gamma(\boldsymbol{\omega}_0)$, i.e., the stopping rule is never satisfied at $t = 0$. Indeed, in the limit as n goes to infinity, with probability one, more that $n^c/2$ iterations are required prior to the stopping rule being triggered. At iteration t in the prior-free asset market mechanism, for $t > n^c/2$ we have

$$d(\boldsymbol{\omega}_t) = \min\left\{n^c, n - 2a_t - 1\right\} = \min\left\{n^c, 2t\right\} = n^c$$

and spacing estimators

$$\sigma^B(\boldsymbol{\omega}_t) = \frac{\theta_{(a_t+1)} - \theta_{(a_t+1+n^c)}}{n^c} \equiv \sigma^B_{a_t+1}$$

and

$$\sigma^S(\boldsymbol{\omega}_t) = \frac{\theta_{[a_t+1+n^c]} - \theta_{[a_t+1]}}{n^c} \equiv \sigma^S_{a_t+1}.$$

Thus, $\phi(\boldsymbol{\omega}_t) = \theta_{(a_t+1)} - (a_t+1)\sigma^B_{a_t+1}$ and $\gamma(\boldsymbol{\omega}_t) = \theta_{[a_t+1]} - (a_t+1)\sigma^S_{a_t+1}$, and the algorithm ends with a_t trades (involving the a_t highest and lowest types) if

$$\theta_{(a_t+1)} - (a_t+1)\sigma^B_{a_t+1} \geq \theta_{[a_t+1]} - (a_t+1)\sigma^S_{a_t+1}.$$

In what follows, we first show that the allocation rule of the prior-free asset market mechanism based on these estimates converges in probability to the allocation rule of the optimal mechanism. Then, second, we show that the ratio of the profit from using the prior-free asset market mechanism to the profit from the optimal mechanism converges in probability to 1 by invoking the payoff-equivalence theorem and continuity to show that the probability that the per-buyer profits and the per-seller expenditures differ from those in the optimum goes to zero as the number of agents goes to infinity.

As above, define p^S and p^B by $1 - F(p^B) = F(p^S)$ and $\Phi(p^B) = \Gamma(p^S)$. By Proposition B.1, in the limit as n goes to infinity, with probability one, the optimal mechanism involves a purchase by every agent with a type greater than p^B and a sale by every agent with a type less than p^S.

Given n, let $\tilde{\Phi}_n(j) \equiv \theta_{(j)} - j\sigma^B_j$ and $\tilde{\Gamma}_n(j) \equiv \theta_{[j]} + j\sigma^S_j$. Let i_n be the random variable such that when there are n agents, $i_n - 1$ agents buy and $i_n - 1$ agents sell in the prior-free asset market mechanism. By the definition of the prior-free asset market mechanism, i_n satisfies

$$\tilde{\Phi}_n(i_n) \geq \tilde{\Gamma}_n(i_n).$$

Using (19) and (20), for all $\varepsilon > 0$,

$$\Pr\left[\sup_{\ell\in\{1,\ldots,n-n^c\}}\left|\tilde{\Phi}_n(\ell)-\Phi(\theta_{(\ell)})\right|>\varepsilon\right]\to 0 \tag{21}$$

and

$$\Pr\left[\sup_{\ell\in\{1,\ldots,n-n^c\}}\left|\tilde{\Gamma}_n(\ell)-\Gamma(\theta_{[\ell]})\right|>\varepsilon\right]\to 0. \tag{22}$$

In other words, as $n\to\infty$, the probability that the thresholds for trade under the prior-free asset market mechanism and under the optimal mechanism differ by more than ε, for any $\varepsilon>0$, goes to zero. This allows us to prove the following lemma:

Lemma B.1. $\operatorname*{plim}_{n\to\infty}\theta_{(i_n)}=p^B$ *and* $\operatorname*{plim}_{n\to\infty}\theta_{[i_n]}=p^S$.

Proof of Lemma B.1. Let $\varepsilon>0$ be given. By the definitions of p^B and p^S,

$$\Phi(p^B-\varepsilon)<\Phi(p^B)=\Gamma(F^{-1}(1-F(p^B)))<\Gamma(F^{-1}(1-F(p^B-\varepsilon))) \tag{23}$$

and

$$\Phi(p^B+\varepsilon)>\Phi(p^B)=\Gamma(F^{-1}(1-F(p^B)))>\Gamma(F^{-1}(1-F(p^B+\varepsilon))). \tag{24}$$

Suppose that

$$\operatorname*{plim}_{n\to\infty}\theta_{(i_n)}<p^B-\varepsilon, \tag{25}$$

which implies that $\operatorname*{plim}_{n\to\infty}\theta_{[i_n]}>F^{-1}(1-F(p^B-\varepsilon))$. Using the continuity of Φ and Γ and Slutsky's Theorem, and the assumption that Φ and Γ are increasing, it follows that

$$\operatorname{plim}\Phi(\theta_{(i_n)})=\Phi(\operatorname{plim}\theta_{(i_n)})<\Phi(p^B-\varepsilon) \tag{26}$$

and

$$\Gamma(F^{-1}(1-F(p^B-\varepsilon)))<\Gamma(\operatorname{plim}\theta_{[i_n]})=\operatorname{plim}\Gamma(\theta_{[i_n]}). \tag{27}$$

By the definition of i_n, $\tilde{\Gamma}(i_n)\le\tilde{\Phi}(i_n)$, so (21)–(22) imply that $\operatorname{plim}\Gamma(\theta_{[i_n]})\le\operatorname{plim}\Phi(\theta_{(i_n)})$. Combining this with (26) and (27), implies that $\Gamma(F^{-1}(1-F(p^B-\varepsilon)))<\Phi(p^B-\varepsilon)$, which contradicts (23), allowing us to conclude that (25) does not hold. Thus, we have $\operatorname*{plim}_{n\to\infty}\theta_{(i_n)}\ge p^B-\varepsilon$. Analogously, using (24), we can conclude that $\operatorname*{plim}_{n\to\infty}\theta_{(i_n)}\le p^B+\varepsilon$, which, letting ε go to zero, completes the proof that $\operatorname*{plim}_{n\to\infty}\theta_{(i_n)}=p^B$. Analogous arguments imply that $\operatorname*{plim}_{n\to\infty}\theta_{[i_n]}=p^S$. ■

Continuation of the proof of Proposition 2. The difference in the per-buyer expected payments between the prior-free asset market mechanism and the optimal mechanism is, by the payoff-equivalence theorem,

$$\left|\int_{\theta_{(i_n)}}^{p^B}\Phi(x)f(x)dx\right|,$$

and similarly the difference in per-seller expected payment is

$$\left|\int_{\theta_{[i_n]}}^{p^S}\Gamma(x)f(x)dx\right|.$$

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

$$\lim_{n\to\infty}\Pr\left(\left|\int_{\theta_{(i_n)}}^{p^B}\Phi(v)f(v)dv\right|>\Delta\right)=0.$$

To show this, note that for any $\varepsilon > 0$,

$$\Pr\left(\left|\int_{\theta_{(i_n)}}^{p^B} \Phi(v) f(v) dv\right| > \Delta\right) \tag{28}$$

$$= \Pr\left(\left|\int_{\theta_{(i_n)}}^{p^B} \Phi(v) f(v) dv\right| > \Delta \mid |\theta_{(i_n)} - p^B| \geq \varepsilon\right) \Pr\left(|\theta_{(i_n)} - p^B| \geq \varepsilon\right) \tag{29}$$

$$+ \Pr\left(\left|\int_{\theta_{(i_n)}}^{p^B} \Phi(v) f(v) dv\right| > \Delta \mid |\theta_{(i_n)} - p^B| < \varepsilon\right) \Pr\left(|\theta_{(i_n)} - p^B| < \varepsilon\right) \tag{30}$$

$$\leq \Pr\left(|\theta_{(i_n)} - p^B| \geq \varepsilon\right) + \Pr\left(\left|\int_{\theta_{(i_n)}}^{p^B} \Phi(v) f(v) dv\right| > \Delta \mid |\theta_{(i_n)} - p^B| < \varepsilon\right). \tag{31}$$

The first term in the last line of the expression above is zero in the limit as n goes to infinity by Lemma B.1, and because $\Phi(v)$ and $f(v)$ are bounded, the second term is also zero in the limit as n goes to infinity because we can make ε arbitrarily small. Thus,

$$\lim_{n\to\infty} \Pr\left(\left|\int_{\theta_{(i_n)}}^{p^B} \Phi(v) f(v) dv\right| > \Delta\right) = 0$$

as required, and similarly for the seller side. ■

Proof of Lemma 1. Suppose to the contrary that the lemma is not true. That is, suppose, for example, that $V > \overline{V}$. By the monotonicity of the virtual type functions, this implies that for all $j \in \mathcal{M}$, $\overline{p}_j^B > p_j^B$ and $\overline{p}_j^S > p_j^S$, and hence $1 - F_j(\overline{p}_j^B) < 1 - F_j(p_j^B)$ and $F_j(\overline{p}_j^S) > F_j(p_j^S)$ for all $j \in \mathcal{M}$, which means that the market clearing condition (12) is violated. An analogous argument applies to the case $V < \underline{V}$, which proves the result. ■

Appendix C. Optimal mechanism details

In this appendix, we provide additional details of the optimal mechanism.

The optimal allocation rule and ironing parameter are defined so that the allocation rule maximizes the expected value of the sum of ironed virtual types, weighted by quantities, and such that the interim expected net trade for agents with types in the ironing region is zero. Adapting Lu and Robert (2001) to our setup, given ironing parameter z, the allocation rule that maximizes $\sum_{i\in\mathcal{N}} V(\theta_i) q_i(\boldsymbol{\theta})$ induces $\frac{n-1}{2}$ trades if half or more of the agents have ironed virtual types below z or above z. In that case, every agent except the median agent trades. Otherwise, the number of trades is determined by the larger of the number of agents with ironed virtual types below z and above z. For example, if some number $s < \frac{n-1}{2}$ of the agents have ironed virtual types less than z and some number $b < \frac{n-1}{2}$ of the agents have ironed virtual types greater than z, and if $s \geq b$, then there are s trades: all the agents with ironed virtual types less than z sell, all the agents with ironed virtual types greater than z buy, and $s - b$ randomly selected agents with ironed virtual types equal to z also buy.

Given this, we can write the optimal number of trades given an odd number n of agents, ironing parameter z, and type realization $\boldsymbol{\theta}$ as

$$\min\left\{\frac{n-1}{2},\ \max\left\{\sum_{i\in\mathcal{N}} \mathbf{1}_{\Gamma(\theta_i)<z}, \sum_{i\in\mathcal{N}} \mathbf{1}_{\Phi(\theta_i)>z}\right\}\right\}.$$

The optimal allocation rule is then defined to reflect these trades, and the definition of the optimal mechanism is completed by choosing z such that for $\theta_i \in [\Gamma^{-1}(z), \Phi^{-1}(z)]$, each agent's interim expected allocation is equal to its initial allocation of 1, i.e., $\hat{q}(\theta_i) \equiv \mathbb{E}_{\boldsymbol{\theta}_{-i}}[q_i(\theta_i, \boldsymbol{\theta}_{-i})] = 1$. As shown below in Lemma C.1, if f is symmetric in the sense that $f(\theta) = f(1-\theta)$ for all $\theta \in [0, 1]$, then the ironing parameter z is such that $z = \Gamma(p^S) = \Phi(p^B)$, where p^S and p^B are as defined in (2) and (3).

Given z and defining p_z^S and p_z^B by $\Gamma(p_z^S) = z = \Phi(p_z^B)$, the optimal allocation rule for agent i is

$$q_i(\theta_i, \boldsymbol{\theta}_{-i}) = \begin{cases} 0 & \text{if } \theta_i < p_z^S \text{ or } \theta_i > p_z^B \text{ and } x < \boldsymbol{\theta}_{-i[\frac{n-1}{2}]}, \\ 1 & \text{if } \theta_i < p_z^S \text{ or } \theta_i > p_z^B \text{ and } \boldsymbol{\theta}_{-i[\frac{n-1}{2}]} < \theta_i < \boldsymbol{\theta}_{-i[\frac{n+1}{2}]}, \\ 2 & \text{if } \theta_i < p_z^S \text{ or } \theta_i > p_z^B \text{ and } \boldsymbol{\theta}_{-i[\frac{n+1}{2}]} < \theta_i, \\ 1 - \frac{b-s}{n-s-b} & \text{if } \theta_i \in [p_z^S, p_z^B],\ s = \sum_{j \neq i} \mathbf{1}_{\theta_j < p_z^S} < \sum_{j \neq i} \mathbf{1}_{\theta_j > p_z^B} = b, \\ 1 & \text{if } \theta_i \in [p_z^S, p_z^B] \text{ and } \sum_{j \neq i} \mathbf{1}_{\theta_j < p_z^S} = \sum_{j \neq i} \mathbf{1}_{\theta_j > p_z^B}, \\ 1 + \frac{s-b}{n-s-b} & \text{if } \theta_i \in [p_z^S, p_z^B],\ s = \sum_{j \neq i} \mathbf{1}_{\theta_j < p_z^S} > \sum_{j \neq i} \mathbf{1}_{\theta_j > p_z^B} = b. \end{cases}$$

This implies that, given z, the interim expected allocation for an agent with an ironed virtual type equal to z does not depend directly on the agent's type and is given by the expectation of the last three lines in the above expression for $q_i(\theta_i, \boldsymbol{\theta}_{-i})$, which we can write as:

$$\begin{aligned} h(z) \equiv & \sum_{b=1}^{n-1} \sum_{s=0}^{\min\{n-b-1,b-1\}} \binom{n-1}{b} \binom{n-1-b}{s} \left(1 - \frac{b-s}{n-s-b}\right) \\ & \cdot F(p_z^S)^s (1 - F(p_z^B))^b (F(p_z^B) - F(p_z^S))^{n-1-s-b} \\ & + \sum_{s=0}^{\frac{n-1}{2}} \binom{n-1}{s} \binom{n-1-s}{s} F(p_z^S)^s (1 - F(p_z^B))^s (F(p_z^B) - F(p_z^S))^{n-1-2s} \\ & + \sum_{s=1}^{n-1} \sum_{b=0}^{\min\{n-s-1,s-1\}} \binom{n-1}{s} \binom{n-1-s}{b} \left(1 + \frac{s-b}{n-s-b}\right) \\ & \cdot F(p_z^S)^s (1 - F(p_z^B))^b (F(p_z^B) - F(p_z^S))^{n-1-s-b}. \end{aligned}$$

Thus, the optimal mechanism requires z such that $h(z) = 1$,[32] giving us the interim expected allocation rule (the same for all agents) of:

$$\hat{q}(\theta) = \begin{cases} \binom{n-1}{\frac{n-1}{2}} F(\theta)^{\frac{n-1}{2}} (1 - F(\theta))^{\frac{n-1}{2}} & \\ +2 \sum_{j=\frac{n+1}{2}}^{n-1} \binom{n-1}{j} F(\theta)^j (1 - F(\theta))^{n-1-j} & \text{if } \theta < p_z^S \text{ or } \theta > p_z^B, \\ 1 & \text{otherwise.} \end{cases}$$

Per-agent expected revenue can be calculated as the expectation of the virtual type times the interim expected number of trades. Noting that each agent starts with one unit, so that agent i's interim expected number of trades is $\hat{q}(\theta_i) - 1$, per-agent expected revenue is

$$\pi^O \equiv \int_0^{p_z^S} \Gamma(x)(\hat{q}(x) - 1) dF(x) + \int_{p_z^B}^1 \Phi(x)(\hat{q}(x) - 1) dF(x).$$

We conclude by providing conditions under which the ironing range in the optimal mechanism is given by $[p^S, p^B]$, where p^S and p^B are as defined in (2) and (3).

Lemma C.1. *Assuming regularity, i.e., Φ and Γ are increasing, and that f is symmetric in the sense that $f(\theta) = f(1-\theta)$ for all $\theta \in [0, 1]$, then the ironing region in the optimal mechanism is $[p^S, p^B]$.*

Proof. Let a_0 and b_0 be the lower and upper bound of the ironing interval. For the interim expected allocation for an agent of type $\theta \in [a_0, b_0]$ to be 1 as required under the pointwise maximizer for the ironed virtual types, it must be that $F(a_0) = 1 - F(b_0)$. Let $p = F^{-1}(1/2)$ be the ironing parameter. The ironed virtual type function with ironing parameter p and assumed ironing interval $[p^S, p^B]$ is such that

[32] Such a z exists because $h(z)$ is continuous, with $p_0^S = 0$ and $p_1^B = 1$, which imply that $h(0) < 1$ and $h(1) > 1$.

$$L = \frac{\int_{p^S}^{p} (\Gamma(x) - p) f(x) dx}{1/2 - F(p^S)} = \frac{\int_{p}^{p^B} (p - \Phi(x)) f(x) dx}{F(p^B) - 1/2} = R.$$

Integrating by parts gives

$$L = \frac{(p - p^S) F(p^S)}{1/2 - F(p^S)} \quad \text{and} \quad R = \frac{(p^B - p)(1 - F(p^B))}{F(p^B) - 1/2}.$$

Using (3) twice, one gets

$$R = \frac{(p^B - p) F(p^S)}{1/2 - F(p^S)}.$$

Thus, $L = R$ is equivalent to $p - p^S = p^B - p$, i.e., to p^B and p^S being symmetric around p, which is equivalent to $p = (p^B + p^S)/2$. But if f is symmetric, then (3) holds if and only if $p^B = 1 - p^S$, which is equivalent to $\frac{p^B + p^S}{2} = 1/2$, which is p. ■

References

Abdulkadiroğlu, A., Angrist, J.D., Narita, Y., Pathak, P., 2017. Research design meets market design: using centralized assignment for impact evaluation. Econometrica 85, 1373–1432.
Abdulkadiroğlu, A., Angrist, J.D., Narita, Y., Pathak, P., 2022. Breaking ties: regression discontinuity design meets market design. Econometrica 90, 117–151.
Agarwal, N., Ashlagi, I., Rees, M., Somaini, P., Waldinger, D., 2021. Equilibrium allocations under alternative waitlist designs: evidence from deceased donor kidneys. Econometrica 89, 37–76.
Baliga, S., Vohra, R., 2003. Market research and market design. Adv. Theor. Econ. 3, 5.
Baranov, O., Aperjis, C., Ausubel, L.M., Morrill, T., 2017. Efficient procurement auctions with increasing returns. Am. Econ. J. Microecon. 9, 1–27.
Börgers, T., 2015. An Introduction to the Theory of Mechanism Design. Oxford University Press.
Bulow, J., Roberts, J., 1989. The simple economics of optimal auctions. J. Polit. Econ. 97, 1060–1090.
Chen, Y.-C., Li, J., 2018. Revisiting the foundations of dominant-strategy mechanisms. J. Econ. Theory 178, 294–317.
Chetty, R., Friedman, J.N., Stepner, M., 2021. Who Spent Their Last Stimulus Checks? NYTimes.
Cramton, P., Gibbons, R., Klemperer, P., 1987. Dissolving a partnership efficiently. Econometrica 55, 615–632.
Cripps, M.W., Swinkels, J.M., 2006. Efficiency of large double auctions. Econometrica 74, 47–92.
Delacrétaz, D., Loertscher, S., Mezzetti, C., 2022. When Walras meets Vickrey. Theor. Econ.
Gresik, T., Satterthwaite, M., 1989. The rate at which a simple market converges to efficiency as the number of traders increases: an asymptotic result for optimal trading mechanisms. J. Econ. Theory 48, 304–332.
Johnson, T.R., 2019. One-Sided Markets. Working Paper. University of Notre Dame.
Krishna, V., 2002. Auction Theory. Elsevier Science, Academic Press.
Laffer, A.B., 2004. The Laffer Curve: Past, Present, and Future. Backgrounder. No. 1765, Heritage Foundation.
Lewis, T.R., Sappington, D.E.M., 1989. Countervailing incentives in agency problems. J. Econ. Theory 49, 294–313.
Leyton-Brown, K., Milgrom, P., Segal, I., 2017. Economics and computer science of a radio spectrum reallocation. Proc. Natl. Acad. Sci. USA 114, 7207–7209.
Li, J., Dworczak, P., 2021. Are Simple Mechanisms Optimal when Agents are Unsophisticated? Working Paper. Northwestern University.
Loertscher, S., Marx, L., Wilkening, T., 2015. A long way coming: designing centralized markets with privately informed buyers and sellers. J. Econ. Lit. 53, 857–897.
Loertscher, S., Marx, L.M., 2020a. Asymptotically optimal prior-free clock auctions. J. Econ. Theory 187, 1–35.
Loertscher, S., Marx, L.M., 2020b. A dominant-strategy asset market mechanism. Games Econ. Behav. 120, 1–15.
Loertscher, S., Marx, L.M., 2022. Bilateral trade with multiunit demand and supply. Manag. Sci.
Loertscher, S., Mezzetti, C., 2021. A dominant strategy, double clock auction with estimation-based tâtonnement. Theor. Econ. 16, 943–978.
Loertscher, S., Wasser, C., 2019. Optimal structure and dissolution of partnerships. Theor. Econ. 14, 1063–1114.
Lu, H., Robert, J., 2001. Optimal trading mechanisms with ex ante unidentified traders. J. Econ. Theory 97, 50–80.
Milgrom, P., 2017. Discovering Prices: Auction Design in Markets with Complex Constraints. Columbia University Press, New York.
Milgrom, P., Segal, I., 2020. Clock auctions and radio spectrum reallocation. J. Polit. Econ. 128, 1–31.
Mussa, M., Rosen, S., 1978. Monopoly and product quality. J. Econ. Theory 18, 301–317.
Myerson, R.B., 1981. Optimal auction design. Math. Oper. Res. 6, 58–73.
Norman, P., 2004. Efficient mechanisms for public goods with use exclusions. Rev. Econ. Stud. 71, 1163–1188.
Segal, I., 2003. Optimal pricing mechanisms with unknown demand. Am. Econ. Rev. 93, 509–529.
Silverman, B.W., 1986. Density Estimation for Statistics and Data Analysis. Chapman and Hall.
Wilson, R., 1987. Game theoretic analysis of trading processes. In: Bewley, T. (Ed.), Advances in Economic Theory. Cambridge University Press, Cambridge, U.K., pp. 33–70.
Wilson, R., 1993. Nonlinear Pricing. Oxford University Press.