Optimal market thickness and clearing*

Simon Loertscher† Ellen V. Muir‡ Peter G. Taylor§

First version: 15 February 2016 This version: December 3, 2018

Abstract

To determine the optimal market clearing policy that balances the benefits of thickness against costly delay, we solve a dynamic mechanism design model with sequential arrival of buyers and sellers with private information about their types. With discrete (binary) types, there is an efficient mechanism that is incentive compatible, individually rational and always balances the budget if the discount factor is sufficiently large (if storing at least one trade is efficient). For large discount factors, most welfare gains from dynamic mechanisms relative to instantaneous clearing are reaped by the simplest dynamic mechanism that clears at an optimally chosen, fixed frequency.

Keywords: market thickness, clearing houses, market mechanisms, two-sided private information, dynamic efficiency, trading venues, order books, (im)possibility of efficient trade.
JEL-Classification: C72, D47, D82

*We thank Eric Budish, Yeon-Koo Che, Darrell Duffie, Alexander Frankel, Dan Garrett, Scott Kominers, Thomas Mariotti, Leslie Marx, Hayden Melton, Claudio Mezzetti, Paul Milgrom, Carlos Oyarzun, Andy Skrzypacz, Alexander Teytelboym, Leeat Yariv and audiences at the University of Melbourne, University of Queensland, Ecole Polytechnique Fédérale de Lausanne, Stanford University, the Melbourne IO and Theory Day, the Econometric Society Australasian Meeting in Sydney 2016, the Econometric Society European Meeting in Geneva 2016, the European Association for Research in Industrial Economics Meeting in Lisbon 2016, the Columbia/Duke/MIT/Northwestern IO Theory Conference in Chicago 2016, the Australasian Economic Theory Workshop in Auckland 2017, the University of North Carolina at Chapel Hill, the University of New South Wales, the Chicago Mercantile Exchange and the University of Chicago for comments and suggestions. Financial support provided by the Australian Research Council through Laureate Fellowship FL130100039, the ARC Centre of Excellence for Mathematical and Statistical Frontiers, the Elizabeth and Vernon Puzey Foundation and the Samuel and June Hordern Endowment is also gratefully acknowledged. Jonatham Lim provided excellent research assistance. We also acknowledge support from the European Research Council under the Grant Agreement no. 340903.

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

‡Department of Economics, Stanford University Email: evmuir@stanford.edu

§School of Mathematics and Statistics, The University of Melbourne, Melbourne, Australia Email: taylorpg@unimelb.edu.au.

1 Introduction

Continuous-time double auctions, which clear compatible trades instantaneously, have proved successful in experimental settings and are widely used in financial markets. Recently, however, the high frequency of trade associated with continuous-time trading mechanisms has come under scrutiny on the ground that they induce socially wasteful arms races into arbitrage technologies. Batch auctions – that is, uniform price auctions run at a fixed frequency – have been proposed as a remedy (Budish et al., 2015). In response to this and other related concerns, a number of major exchanges have abandoned the use of continuous-time market mechanisms. For example, IEX, a US-based stock exchange, incorporated a delay of 350 microseconds into its mechanism to eliminate the speed advantages required for certain predatory trading strategies (Lewis, 2014). Another case in point is Thomson Reuters Matching, a major interbank electronic trading venue in the foreign exchange market, which abandoned its continuous-time market mechanism by introducing a buffer time to de-emphasize speed, with the buffer being triggered by trading behavior (Melton, 2017). Similar "speedbump" trading mechanisms are currently under consideration at both the New York Stock Exchange and the Chicago Stock Exchange (McCrank, 2017b,a).

While these developments show increased awareness of the importance of how markets are cleared, they also highlight how little is known in general about the basic question of how, and how often, markets should be cleared. Accumulating traders increases market thickness, which is good insofar as it offers additional or more valuable opportunities to trade, but bad to the extent that the reduced speed creates costly delay. To date, practitioners have received little to no guidance from economics about the optimal design of market clearing mechanisms and have, perhaps as a consequence, paid little attention to the economic tradeoffs that are involved. For example, the transition from paper to computer-organized trading at the New York Stock Exchange was exclusively driven by the programmer's desire to execute trades as fast as possible without any consideration of the tradeoff between speed and market thickness. Likewise, the Native Vegetation Exchange (NVX) for Victoria, Australia, was designed to execute compatible trades instantaneously, not on the grounds that this would be optimal but because of computational complexity. Similarly, eBay's clearing mechanism does not allow traders to accumulate, which, as documented by Hendricks and Sorensen (2018), results in substantial welfare losses.

In this paper, we take the first step towards closing this gap. We derive the optimal market clearing mechanism in a model with sequential arrival of buyers and sellers who are privately informed about their values and costs. We show that the impossibility of efficient, incentive compatible and individually rational trade with privately informed agents is overcome provided the discount factor is sufficiently large. With binary types this occurs as soon as it is optimal to store at least one trade. To the best of our knowledge, our paper thus brings to light an important novel aspect to the debates pertaining to the efficiency of secondary markets going back to, at least, Lerner (1944), Coase (1960), Vickrey (1961), Hurwicz (1972), and Myerson and Satterthwaite (1983). While our results have a Coasian flavour in that they show that the inefficiency of initial allocations can be resolved with an appropriately designed dynamic market clearing mechanism, they also provide a strong basis, and demonstrate the need, for market design: without the dynamic mechanism, inefficiencies from the initial allocation cannot always be resolved.[1,2] Furthermore, with binary types we show that efficient trade is possible precisely when the efficient allocation can be implemented using a posted price mechanism. Because of their simplicity, such mechanisms are also independently appealing for practical purposes.[3]

For large discount factors and general discrete types we show that most welfare gains from using dynamic mechanisms relative to instantaneous clearing are reaped by the simplest

[1]Both a similarity with and difference to Cramton et al. (1987) are worth noting here. Recall that Cramton et al. show that in a static setup efficient reallocation via a centralized market mechanism is possible when the initial ownership structure is sufficiently symmetric. We show that with an appropriately designed market mechanism efficient reallocation is possible even with extreme initial ownership, provided only impatience is not too severe. Common to both papers is that they make strong cases, implicitly or explicitly, for market design. Neither in the setting of Cramton et al. nor in ours will non-intermediated bilateral trade lead to efficiency in any degree of generality.

[2]There is an important difference between our possibility result and the recent literature on the (im)possibility of efficient bilateral trade in repeated settings that started with Athey and Miller (2007); see also the literature review below and Garrett (2016) for further references. In repeated settings, the efficient policy does not vary with the discount factor. More patience merely means that the individual rationality constraints become more slack. In our setting, in contrast, it is precisely the change in the efficient policy resulting from increases in the discount factor that renders efficient trade without a deficit possible.

[3]There is an interesting analogy to static settings. For the static bilateral trade setup with overlapping supports, the impossibility theorem of Myerson and Satterthwaite (1983) holds. Consequently, posted prices constrain social surplus while avoiding a deficit (Hagerty and Rogerson, 1987). If the supports do not overlap, efficient trade is possible and can be implemented with posted prices. In the dynamic setting, storing one trade has thus the same effects as have non-overlapping supports in the static setup.

dynamic mechanism that clears at an optimally chosen, fixed frequency. We also show that, for a sufficiently large discount factor, a profit-targeting exchange generates greater social welfare gains than a welfare-targeting market maker that uses a less sophisticated form of market clearing. This suggests that traders may prefer to trade via a large monopolist exchange rather than via a periodic ex post efficient exchange that never stores any trades. Moreover, we show that an ad valorem tax imposed on the profit of a profit-maximizing market maker does not distort the market maker's policy whereas a specific tax does. Thus, our paper also sheds new light on the effects of different forms of transaction taxes which featured prominently in policy debates following the Global Financial Crisis and are relevant to other regulatory issues, for example artist remuneration for online content streaming[4].

This paper relates, first and foremost, to the literature on mechanism and market design. In particular, we apply the techniques developed by Myerson (1981) to a dynamic setting with discrete types and two-sided private information. Static versions of this setup with two-sided private information have previously been studied by, among others, Myerson and Satterthwaite (1983), Baliga and Vohra (2003), and Loertscher and Marx (2017). We use the notions of interim and period ex post incentive compatibility that were introduced and first used by Bergemann and Välimäki (2010). Much of the recent literature on dynamic mechanism design, including Athey and Miller (2007), Bergemann and Välimäki (2010), Athey and Segal (2013), Pavan et al. (2014) and Skrzypacz and Toikka (2015), considers settings in which a static population of agents receives private information over time. In contrast, our paper considers a dynamic population of agents with persistent types. In such setups, the current allocation decision determines the set of feasible allocations in future periods and the designer faces the optimal timing problem of deciding when to run a static mechanism. Recent contributions to this strand of literature include Parkes and Singh (2003), Gershkov and Moldovanu (2010) and Board and Skrzypacz (2016). However, none of the aforementioned papers explicitly address the optimal timing problem[5] nor do they consider varying degrees of sophistication of the mechanisms or compare welfare and profit

[4]In particular, some have advocated for artists being paid a fixed amount per song played and others favouring a share of the streaming service's revenue proportional to the time an artist's songs are played.

[5]Recent papers that address the optimal timing problem in one-sided settings include Pai and Vohra (2013), Mierendorff (2013) and Mierendorff (2016).

maximization.[6]

Our paper is methodologically related to the recent literature on dynamic matching where monetary transfers cannot be used to incentivize agents. For example, building on Ünver (2010) and Anderson et al. (2017), Akbarpour et al. (2017) study efficiency in a dynamic matching model with complete information in which exchange possibilities have a network structure.[7] Most importantly, our paper draws inspiration from the work of Baccara et al. (2016). Motivated by the problem of matching children and parents in an adoption "market" they consider a dynamic, two-sided matching problem. The efficient algorithm Baccara et al. derive is similar to the optimal market clearing policy in our paper when we specialize the setup to binary types. There are, however, crucial differences between our approach and that of Baccara et al. (2016). We adhere to the standard assumption in the dynamic mechanism design literature of geometric discounting[8], whereas Baccara et al. assume each agent incurs a fixed per period cost of delay. This captures a notion of agents' aging, reducing and eventually eliminating the gains from "trade", which seems appropriate for the application at hand. In contrast, our assumptions of geometric discounting and quasilinear payoffs permit the use of monetary transfers to incentivize agents and allow us to study a broad range of questions that Baccara et al. cannot address such as revenue maximization by the market maker, the possibility of efficient trade without running a deficit, or implementation via price posting and the equilibrium distribution of prices. Thus, while some elements such as the efficient allocation rule with binary types are naturally similar, the analyzes Baccara et al. and we perform shed light on fundamentally different questions and thereby complement each other.

Our paper also relates to the fundamental question about the (im)possibility of efficient trade. Coase (1960) made the important point that policy debates about the (initial) allocation of property rights necessarily center around the question of transaction costs. Vickrey (1961), Hurwicz (1972) and Myerson and Satterthwaite (1983) argued forcefully that private

[6] There is also a vast literature on intermediation in financial markets; see, for example, Mendelson (1982) or Kelly and Yudovina (2016) and references therein.

[7] For a recent paper that exploits a connection between static mechanism design models and matching *with* transfers, see Delacrétaz et al. (2018).

[8] See, for example, Athey and Miller (2007), Bergemann and Välimäki (2010), Athey and Segal (2013), Pavan et al. (2014) and Skrzypacz and Toikka (2015).

information can be an insurmountable transaction cost while Cramton et al. (1987) pointed out that the answer to the question of whether efficient trade is possible depends on the initial allocation of property rights. Milgrom (2017) provides persuasive arguments that complexity may be an additional source of transaction costs impeding efficient (re)allocation.[9] Our paper contributes to this debate by pointing to the importance of dynamic aspects. In particular, when impatience is not too large, we show that efficiency is possible without running a deficit. Yet, there is considerable scope for market design as the institutions that enable efficiency may not arise spontaneously through quick trial and error processes.

The remainder of this paper is organized as follows. Section 2 introduces the model and key concepts. In Section 3, we solve the mechanism design problem for a symmetric setup with binary types and discuss key properties of the optimal mechanism. Section 4 analyzes posted price mechanisms and relates these, among other things, via the efficient market clearing policy to the (im)possibility of efficient trade. Section 5 compares the performance of the optimal dynamic mechanism to two less sophisticated dynamic mechanisms and instantaneous market clearing. Section 6 concludes. All proofs, generalizations and algorithms are provided in the Appendix.

2 Model

In this section, we introduce the general setup, the designer's mechanism design problem and some basic results that follow from applying mechanism design theory.

2.1 General setup

We first introduce traders' types and payoff functions and define the arrival process. We then introduce the objective of the designer and define the mechanism design problem faced by the designer.

We consider a discrete-time infinite horizon setup in which a market designer operates a two-sided exchange. In each period $t \in \mathbb{N}$ a single buyer B_t and a single seller S_t arrive (see

[9]There is also empirical evidence suggesting that private information and the no deficit constraint are not the only sources of transaction costs. For example, Larsen (2018) shows that the bargaining outcomes in wholesale used-car auctions are well below the second-best frontier.

Online Appendix D for generalizations of the arrival process).[10] All agents and the designer are risk neutral geometric discounters, with a common discount factor $\delta \in [0, 1)$. We assume that all agents have quasilinear preferences and that each buyer demands at most one unit and each seller has the capacity to produce at most one unit. We also assume that agents can only trade via the designer's platform and the value of agents' outside option of not participating is zero.

Assume that buyers draw their types independently from a discrete distribution F with probability mass function f whose support is given by $\mathcal{V} := \{v_1, \ldots, v_n\}$ with $v_1 < \cdots < v_n$ for $n \in \mathbb{N}$ and that sellers draw their types independently from a discrete distribution G with probability mass function g and support $\mathcal{C} := \{c_1, \ldots, c_m\}$, where $c_1 < \cdots < c_m$ for $m \in \mathbb{N}$. We will occasionally refer to v_1 and c_m as the least-efficient buyer and seller type, respectively. The arrival process, δ and the distributions F and G are common knowledge and arrivals are observable.

This setup has a number of advantages. Private information about values and costs makes the price discovery problem associated with market making in two-sided settings non-trivial. The assumption of independently distributed private types implies that the optimal Bayesian mechanism provides a practical benchmark. For static settings, it is well-known that under these assumptions inducing efficient trade by privately informed buyers and sellers is not possible without running a deficit if, in addition, values and costs are distributed according to absolutely continuous distribution functions with identical and compact supports.[11] Besides its obvious real-world appeal, the setting with private information has the benefit that it neither presumes nor precludes efficiency.[12] It also makes the problem of revenue generation interesting in a plausible way. We depart from the standard Myersonian setup by assuming

[10]The assumption that arrival is pairwise is made only to simplify the exposition. It is almost without loss of generality as we show in detail in Online Appendix D.

[11]With correlated types, full surplus extraction and efficiency are possible using mechanisms à la Crémer and McLean (1985, 1988), which are not very robust in a number of relevant dimensions including wealth constraints; see, for example, Kosmopoulou and Williams (1998), and Börgers (2015). With interdependent types, depending on fine details such as distributional assumptions, efficiency may be possible using more elaborate mechanisms than direct, one-shot revelation mechanisms (Mezzetti, 2004).

[12]That private information is at times an insurmountable transaction cost has long been recognized; see, for example, Vickrey (1961), Hurwicz (1972) and Myerson and Satterthwaite (1983). Therefore, private information is one way of avoiding the "Coasian Irrelevance" (Che, 2006) associated with the Coase Theorem (Coase, 1960; Stigler, 1966).

discrete types to make the state space and model tractable. We will discuss how discreteness affects the results as we go. Finally, geometric discounting is the natural assumption for dynamic mechanism design settings in which the designer has to incentivize agents to reveal information and cares about revenue.

2.2 The mechanism design problem

The designer's problem is to find an incentive compatible, individually rational mechanism that maximizes her objective. Denoting by $\langle \boldsymbol{Q}, \boldsymbol{M} \rangle$ a direct, feasible mechanism that is incentive compatible and individually rational in ways that will be explained shortly, we let $R(\langle \boldsymbol{Q}, \boldsymbol{M} \rangle)$ and $W(\langle \boldsymbol{Q}, \boldsymbol{M} \rangle)$ denote, respectively, the expected discounted revenue and social welfare gain generated by the mechanism. In the tradition of Myerson and Satterthwaite (1983) and Gresik and Satterthwaite (1989), we assume that the designer is interested in constrained efficient mechanisms. These mechanisms maximize W subject to the constraint of generating an expected revenue of at least $\underline{R}$,[13] as well as the appropriate incentive compatibility and individual rationality constraints. It is well known that the set of constrained efficient mechanisms is the set of Bayesian optimal mechanisms that, for any given $\alpha \in [0, 1]$, maximize the Ramsey objective

$$\alpha R(\langle \boldsymbol{Q}, \boldsymbol{M} \rangle) + (1 - \alpha) W(\langle \boldsymbol{Q}, \boldsymbol{M} \rangle), \tag{1}$$

where the maximum is taken over feasible, incentive compatible and individually rational mechanisms. Notice that when $\alpha = 0$ and $\alpha = 1$, we obtain, respectively, the efficient and profit-maximizing mechanism.

Incentive compatibility and individual rationality

We restrict ourselves to direct, truthful mechanisms. A direct mechanism $\langle \boldsymbol{Q}, \boldsymbol{M} \rangle$ consists of an allocation rule $\boldsymbol{Q} = \{\boldsymbol{Q}_t\}_{t \in \mathbb{N}}$ and a payment rule $\boldsymbol{M} = \{\boldsymbol{M}_t\}_{t \in \mathbb{N}}$. Let $\mathcal{H}^t := (\mathcal{V} \times \mathcal{C})^t$ be the set of histories of agents' reports up to and including period t. The period t allocation rule $\boldsymbol{Q}_t : \mathcal{H}^t \to \{0, 1\}^{2t}$ maps the period t history of agent reports $\boldsymbol{h}_t$ to the set of period t allocations, and similarly, the period t transfer rule $\boldsymbol{M}_t : \mathcal{H}_t \to \mathbb{R}^{2t}$ maps this history to

[13] Or, equivalently, maximize R subject to the constraint of generating expected welfare gains of at least $\underline{W}$.

the set of period t transfers. Because of the revelation principle, the restriction to direct mechanisms is without loss of generality.

The Bayesian incentive compatibility constraints require that truthful reporting is a best response for every agent, assuming that all other agents report truthfully. The interim individual rationality constraints require that agents' interim expected payoffs are non-negative. Formally, for a given direct mechanism $\langle \boldsymbol{Q}, \boldsymbol{M} \rangle$, let $q(\hat{\theta})$ denote the discounted probability of trade for an agent that arrives in a given period and reports $\hat{\theta}$. Similarly, let $m(\hat{\theta})$ denote the expected payment made (received) by a buyer (seller) that arrives in this period and reports $\hat{\theta}$. Then the *Bayesian incentive compatibility (BIC)* and *individual rationality (IR)* constraints require that for all $v \in \mathcal{V}$ and $c \in \mathcal{C}$,

$$v = \arg\max_{\hat{\theta} \in \mathcal{V}} \left\{ vq(\hat{\theta}) - m(\hat{\theta}) \right\} \quad \text{and} \quad c = \arg\max_{\hat{\theta} \in \mathcal{C}} \left\{ m(\hat{\theta}) - cq(\hat{\theta}) \right\}, \tag{BIC}$$

and

$$vq(v) - m(v) \geq 0 \quad \text{and} \quad m(c) - cq(c) \geq 0. \tag{IR}$$

BIC is equivalent to requiring that the expected discounted allocation for buyers increases in their report and the expected discounted allocation for sellers decreases in their report. For any $\alpha > 0$,[14] it is well known that individual rationality constraints bind for the least-efficient types and that, for all other types, the incentive compatibility constraints bind locally downward for buyers and locally upward for sellers; see, for example, Elkind (2007). Below, when we speak of binding incentive compatibility and individual rationality constraints, we will mean that the incentive compatibility constraints are locally downwards (upwards) binding for buyers (sellers) and the individual rationality constraints are binding for the least-efficient types.

In our setting, as we will show, alternative, stronger notions of incentive compatibility that have been discussed in the literature are equivalent to (BIC).[15] The *interim incentive compatibility constraints (i-IC)* (see, for example, Bergemann and Välimäki, 2010) require

[14]When $\alpha = 0$, there is an indeterminacy: whether the individual rationality constraints bind does not affect the objective and because of the discrete type space, the allocation rule does not pin down payments. By treating $\alpha = 0$ as the limit of $\alpha \to 0$, this indeterminacy can be avoided because it implies that for a given allocation rule, the incentive compatible transfers are revenue maximizing, which in turn implies that the individual rationality constraints bind.

[15]That is, equivalent from the ex ante perspective of the designer; see Proposition 1.

that truthful reporting is optimal for every period t agent and every history $\boldsymbol{h}_{t-1} \in \mathcal{H}_{t-1}$, assuming all other agents report truthfully. Formally, let $q(\hat{\theta}, \boldsymbol{h}_{t-1})$ and $m(\hat{\theta}, \boldsymbol{h}_{t-1})$ denote the discounted probability of trade and expected discounted payment, respectively, for an agent that reports $\hat{\theta}$ at history $\boldsymbol{h}_{t-1}$. For every history $\boldsymbol{h}_{t-1} \in \mathcal{H}_{t-1}$, $v \in \mathcal{V}$ and $c \in \mathcal{C}$, (i-IC) then requires

$$\begin{aligned} v &= \arg\max_{\hat{\theta} \in \mathcal{V}} \left\{ vq(\hat{\theta}, \boldsymbol{h}_{t-1}) - m(\hat{\theta}, \boldsymbol{h}_{t-1}) \right\}, \\ c &= \arg\max_{\hat{\theta} \in \mathcal{C}} \left\{ m(\hat{\theta}, \boldsymbol{h}_{t-1}) - cq(\hat{\theta}, \boldsymbol{h}_{t-1}) \right\}. \end{aligned} \tag{i-IC}$$

Similarly, *(periodic) interim individual rationality constraints (i-IR)* require that, for every history $\boldsymbol{h}_{t-1} \in \mathcal{H}_{t-1}$, $v \in \mathcal{V}$ and $c \in \mathcal{C}$,

$$vq(v, \boldsymbol{h}_{t-1}) - m(v, \boldsymbol{h}_{t-1}) \geq 0 \quad \text{and} \quad m(c, \boldsymbol{h}_{t-1}) - cq(c, \boldsymbol{h}_{t-1}) \geq 0. \tag{i-IR}$$

Finally, *periodic ex post incentive compatibility constraints (P-IC)* require that truthful reporting is optimal for every period t agent and every history $\boldsymbol{h}_{t-1}$, regardless of the report of the other period t agent, assuming that all other agents report truthfully. Formally, let $q(\hat{\theta}, \vartheta, \boldsymbol{h}_{t-1})$ and $m(\hat{\theta}, \vartheta, \boldsymbol{h}_{t-1})$ denote the discounted probability of trade and expected discounted payment, respectively, for an agent that reports $\hat{\theta}$ at history $\boldsymbol{h}_{t-1}$ when the other period t agent reports ϑ. For every history $\boldsymbol{h}_{t-1} \in \mathcal{H}_{t-1}$, $v \in \mathcal{V}$ and $c \in \mathcal{C}$, (P-IC) requires

$$\begin{aligned} v &= \arg\max_{\hat{\theta} \in \mathcal{V}} \left\{ vq(\hat{\theta}, c, \boldsymbol{h}_{t-1}) - m(\hat{\theta}, c, \boldsymbol{h}_{t-1}) \right\}, \\ c &= \arg\max_{\hat{\theta} \in \mathcal{C}} \left\{ m(\hat{\theta}, v, \boldsymbol{h}_{t-1}) - cq(\hat{\theta}, v, \boldsymbol{h}_{t-1}) \right\}. \end{aligned} \tag{P-IC}$$

Similarly, *periodic ex post individual rationality constraints (P-IR)* require that, for every history $\boldsymbol{h}_{t-1} \in \mathcal{H}_{t-1}$, $v \in \mathcal{V}$ and $c \in \mathcal{C}$,

$$vq(v, c, \boldsymbol{h}_{t-1}) - m(v, c, \boldsymbol{h}_{t-1}) \geq 0 \quad \text{and} \quad m(c, v, \boldsymbol{h}_{t-1}) - cq(c, v, \boldsymbol{h}_{t-1}) \geq 0. \tag{P-IR}$$

Feasibility Beyond individual rationality and incentive compatibility constraints, the designer also has to respect physical feasibility constraints. While the problem of eliciting information about agents' types truthfully is a static problem because each agent has a time-invariant type which only needs to be elicited upon the agent's arrival, when dealing with the dynamic aspects of the mechanism design problem, we need to distinguish agents not only by their types but also by their time of arrival, which (because of the pairwise arrival) coincides with their identity. With that in mind, given a period t history $\boldsymbol{h}_t$, we denote the respective period t allocations of buyer and seller $i \in \{1, \ldots, t\}$ as $Q_t^{B_i}(\boldsymbol{h}_t)$ and $Q_t^{S_i}(\boldsymbol{h}_t)$. Similarly, $M_t^{B_i}(\boldsymbol{h}_t)$ and $M_t^{S_i}(\boldsymbol{h}_t)$ denote the respective expected payments from B_i and to S_i in period t given $\boldsymbol{h}_t$.

Feasibility requires that, for all $t \in \mathbb{N}$ and all $\boldsymbol{h}_t \in \mathcal{H}_t$,

$$\sum_{i=1}^{t} Q_t^{B_i}(\boldsymbol{h}_t) \leq \sum_{i=1}^{t} Q_t^{S_i}(\boldsymbol{h}_t) \tag{2}$$

and, for all $i \in \{1, \ldots, t\}$,

$$\sum_{j=i}^{t} Q_j^{B_i}(\boldsymbol{h}_j) \leq 1 \quad \text{and} \quad \sum_{j=i}^{t} Q_j^{S_i}(\boldsymbol{h}_j) \leq 1. \tag{3}$$

Of course, (2) will hold with equality under an optimal mechanism.[16]

2.3 Mechanism design results

For $i \in \{1, \ldots, n-1\}$ and $j \in \{2, \ldots, m\}$, the virtual type functions are

$$\Phi(v_i) = v_i - (v_{i+1} - v_i)\frac{1 - F(v_i)}{f(v_i)} \quad \text{and} \quad \Gamma(c_j) = c_j + (c_j - c_{j-1})\frac{G(c_{j-1})}{g(c_j)}, \tag{4}$$

while for $i = n$ and $j = 1$, they are $\Phi(v_n) = v_n$ and $\Gamma(c_1) = c_1$. As pointed out by Bulow and Roberts (1989), virtual values (virtual costs) have the interpretation of marginal revenue (marginal cost) once one accounts for the agents' private information, treating the probability of trade as the quantity demanded (supplied).

We impose the regularity condition of Myerson (1981). That is, we require that Φ and Γ are increasing. Some of our comparative statics results rely on a stronger *dynamic regularity* condition. We say distributions F and G satisfy *dynamic regularity* if Φ and Γ are non-decreasing and if, for $i \in \{1, \ldots, n-1\}$ and $j \in \{2, \ldots, m\}$,

$$\Phi(v_{i+1}) - \Phi(v_i) > v_{i+1} - v_i \quad \text{and} \quad \Gamma(c_j) - \Gamma(c_{j-1}) > c_j - c_{j-1} \tag{5}$$

[16] Observe that the feasibility constraints captured by (3) are exactly the same as those in a standard assignment game. The additional constraint (2) accounts for the dynamic nature of the problem by making sure that at no point in time aggregate demand exceeds aggregate supply. Note also that the irreversibility of time implies that for any two histories $\boldsymbol{h}_t$ and $\boldsymbol{h}'_s$ with $t < s$ such that $\boldsymbol{h}_t = \boldsymbol{h}'_s$, we have, for all $i, j \leq t$, $Q_j^{B_i}(\boldsymbol{h}_j) = Q_j^{B_i}(\boldsymbol{h}'_j)$ and $Q_j^{S_i}(\boldsymbol{h}_j) = Q_j^{S_i}(\boldsymbol{h}'_j)$.

hold.[17]

For $i \leq t$ let $v^{B_i}(\boldsymbol{h}_t) \in \mathcal{V}$ and $c^{S_i}(\boldsymbol{h}_t) \in \mathcal{C}$ denote the types of buyer B_i and seller S_i given history $\boldsymbol{h}_t \in \mathcal{H}_t$, respectively. Expected discounted social welfare gains under any direct, truthful mechanism that implements the allocation rule $\boldsymbol{Q}$ is then given by

$$W(\boldsymbol{Q}) = \sum_{t=1}^{\infty} \sum_{i=1}^{t} \sum_{\boldsymbol{h}_t \in \mathcal{H}_t} \delta^{t-1} \left(v^{B_i}(\boldsymbol{h}_t) Q_t^{B_i}(\boldsymbol{h}_t) - c^{S_i}(\boldsymbol{h}_t) Q_t^{S_i}(\boldsymbol{h}_t) \right) \mathbb{P}(\boldsymbol{H}_t = \boldsymbol{h}_t). \tag{6}$$

The designer's expected discounted profit under binding incentive compatibility and individual rationality can be determined and expressed in terms of virtual types.

Proposition 1. *Expected discounted profit under any direct mechanism with allocation rule* $\boldsymbol{Q}$ *and (BIC) and (IR) binding is given by*

$$R(\boldsymbol{Q}) = \sum_{t=1}^{\infty} \sum_{i=1}^{t} \sum_{\boldsymbol{h}_t \in \mathcal{H}_t} \delta^{t-1} \left(\Phi(v^{B_i}(\boldsymbol{h}_t)) Q_t^{B_i}(\boldsymbol{h}_t) - \Gamma(c^{S_i}(\boldsymbol{h}_t)) Q_t^{S_i}(\boldsymbol{h}_t) \right) \mathbb{P}(\boldsymbol{H}_t = \boldsymbol{h}_t). \tag{7}$$

Furthermore, $R(\boldsymbol{Q})$ *given in (7) is also the expected discounted profit under any direct mechanism with allocation rule* $\boldsymbol{Q}$ *and (i-IC) and (i-IR) or (P-IC) and (P-IR) binding.*

An immediate implication of this proposition is that the designer cannot increase her payoff by concealing the history $\boldsymbol{h}_{t-1} \in \mathcal{H}_{t-1}$ from arriving period t agent. Using (6) and (7), we can now rewrite the Ramsey objective (1), incorporating incentive compatibility and individual rationality constraints, as

$$\begin{aligned} &\alpha R(\boldsymbol{Q}) + (1-\alpha) W(\boldsymbol{Q}) \\ = &\sum_{t=1}^{\infty} \sum_{i=1}^{t} \sum_{\boldsymbol{h}_t \in \mathcal{H}_t} \delta^{t-1} \left(\Phi_\alpha(v^{B_i}(\boldsymbol{h}_t)) Q_t^{B_i}(\boldsymbol{h}_t) - \Gamma_\alpha(c^{S_i}(\boldsymbol{h}_t)) Q_t^{S_i}(\boldsymbol{h}_t) \right) \mathbb{P}(\boldsymbol{H}_t = \boldsymbol{h}_t), \end{aligned} \tag{8}$$

where, for $\alpha \in [0,1]$, $v \in \mathcal{V}$ and $c \in \mathcal{C}$,

$$\Phi_\alpha(v) := (1-\alpha)v + \alpha\Phi(v) \quad \text{and} \quad \Gamma_\alpha(c) := (1-\alpha)c + \alpha\Gamma(c) \tag{9}$$

are the weighted virtual types. The designer's problem is now to determine the allocation rule $\boldsymbol{Q}_\alpha$ that maximizes (8), subject to the appropriate incentive constraints. Using terminology

[17] For example, uniform distributions satisfy dynamic regularity. With continuous distributions, dynamic regularity simply amounts to assuming that the hazard rate $f(v)/(1-F(v))$ is increasing.

that is standard in mechanism design, we refer to the allocation rule $\boldsymbol{Q}_0$ and the mechanism that implements it as *efficient*. For $\alpha > 0$, we refer to the allocation rule $\boldsymbol{Q}_\alpha$ and the corresponding mechanism as *optimal*.[18]

2.4 Symmetric binary types setup

For the remainder of the main body of the paper, with the exception of Proposition 5 in Section 4, we specialize the setup to one with binary types, and symmetric distributions. That is, we now assume $\mathcal{V} = \{\underline{v}, \overline{v}\}$ and $\mathcal{C} = \{\underline{c}, \overline{c}\}$, normalize $\overline{v} = 1$, $\underline{v} = \Delta_0$, $\underline{c} = 0$, and $\overline{c} = 1 - \Delta_0$ with $\Delta_0 \in (0, 1/2)$, and we impose symmetric distributions by assuming that $\Pr(c = \underline{c}) = p = \Pr(v = \overline{v})$. We refer to this as the *symmetric binary type setting*, with symmetry pertaining to the type structure, the distributions, and the arrival. We refer to buyers of type $\overline{v}$ and sellers of type $\underline{c}$ as *efficient* (and to buyers of type $\underline{v}$ and sellers of type $\overline{c}$ as *inefficient*).

This symmetric binary setup simplifies considerably the exposition, and analysis. As it turns out, little is lost in terms of general insight. We discuss along the way which results generalize to the setup from Section 2.1 and in Online Appendix B we provide the relevant generalizations. Further extensions can be found in Online Appendix D.

As $\delta \to 1$, the setup collapses to a static environment in which there is a continuum of traders and p is the proportion of sellers of type $\overline{v}$ and buyers of type $\underline{c}$. The Walrasian equilibrium for this setup is illustrated in Figure 1. Our assumptions make sure that in the static setup with a continuum of agents efficient buyers and sellers trade in the Walrasian market and inefficient buyers and sellers remain inactive. These assumptions also imply that bilateral trade between a high-value buyer and a high-cost seller (or a low-value buyer and a low-cost seller) generates positive social surplus. Upon the arrival of a $(\overline{v}, \overline{c})$ pair or a $(\underline{v}, \underline{c})$ pair, this induces the trade-off between reaping the (small) gains from trade now and waiting in the hope of creating larger gains from trade in the future.

Note that, with binary types and symmetric distributions, the virtual types of the inef-

[18]For example, with private values the Vickrey (or second-price) auction is an efficient auction. In contrast, the selling mechanism Myerson (1981) derived is called an optimal auction.

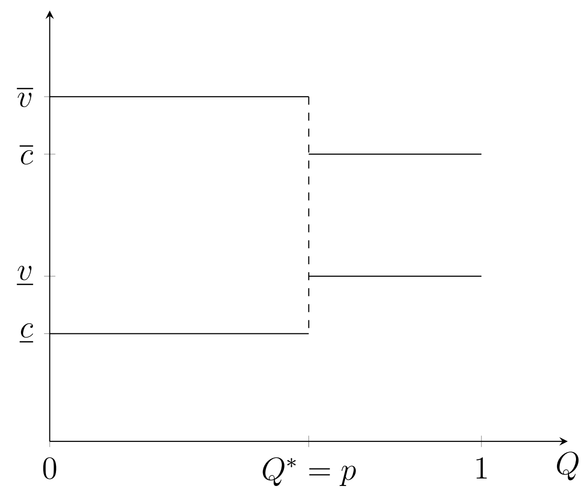

Figure 1: In every period, a buyer-seller pair arrives. Buyers and sellers draw their values and costs independently from the distributions $\{\underline{v}, \overline{v}\}$ and $\{\underline{c}, \overline{c}\}$, respectively, with $\overline{v} > \overline{c} > \underline{v} > \underline{c}$, probability p on $\overline{v}$ and on $\underline{c}$ and the common discount factor δ. For $\delta = 1$, p is the Walrasian quantity.

ficient types become

$$\Phi(\underline{v}) := \underline{v} - \frac{p}{1-p}(\overline{v} - \underline{v}) \quad \text{and} \quad \Gamma(\overline{c}) := \overline{c} + \frac{p}{1-p}(\overline{c} - \underline{c}), \tag{10}$$

while $\Phi(\overline{v}) = \overline{v}$ and $\Gamma(\underline{c}) = \underline{c}$.[19]

Given any $\alpha \in [0,1]$, we let $\Delta_\alpha := \Phi_\alpha(\overline{v}) - \Gamma_\alpha(\overline{c}) = \Phi_\alpha(\underline{v}) - \Gamma_\alpha(\underline{c})$. Observe that

$$\Delta_\alpha = \Delta_0 - \alpha \frac{p}{1-p}(1 - \Delta_0) \leq \Delta_0,$$

where the inequality is strict for $\alpha > 0$.

With binary types, binding individual rationality and incentive compatibility constraints implies $m(\underline{v}) = \underline{v}q(\underline{v})$, $m(\overline{c}) = \overline{c}q(\overline{c})$, $\overline{v}q(\overline{v}) - m(\overline{v}) = \overline{v}q(\underline{v}) - m(\underline{v})$ and $m(\underline{c}) - \underline{c}q(\underline{c}) = m(\overline{c}) - \underline{c}q(\overline{c})$, giving

$$m(\overline{v}) = \overline{v}(q(\overline{v}) - q(\underline{v})) + \underline{v}q(\underline{v}) \quad \text{and} \quad m(\underline{c}) = \underline{c}(q(\underline{c}) - q(\overline{c})) + \overline{c}q(\overline{c}). \tag{11}$$

The incentive compatibility constraints for the worst-off types are satisfied if and only if $q(\overline{v}) \geq q(\underline{v})$ and $q(\underline{c}) \geq q(\overline{c})$.

[19]Observe also that, with binary types, $\Phi(\overline{v}) > \Phi(\underline{v})$ and $\Gamma(\underline{c}) < \Gamma(\overline{c})$ is always the case, eliminating any need or scope to impose monotonicity of virtual type functions separately.

3 Optimal mechanisms in the symmetric binary setup

We now derive the allocation rule $\boldsymbol{Q}_\alpha$ that point by point maximizes (8) as implied by incentive compatibility and individual rationality, temporarily neglecting the constraint that this rule be incentive compatible. Then we verify that the pointwise maximizer permits incentive compatible implementation. In contrast to standard, static mechanism design settings, where the pointwise maximizer is typically trivial, in our setting substantial work goes into its derivation.

3.1 The optimal allocation rule

We begin with two elementary but useful observations. First, notice that $\Delta_\alpha < 0$ is equivalent to the designer wanting to induce trade only between efficient buyers and sellers. Consequently, for $\Delta_\alpha < 0$, the optimal allocation rule induces trade if and only if the buyer has value $\overline{v}$ and a the seller's cost is $\underline{c}$.[20] For remainder of the derivation of the optimal allocation rule, we therefore assume that the parameters α, p and Δ_0 are such that

$$\Delta_\alpha > 0. \tag{12}$$

Since $\Delta_\alpha \leq \Delta_0 < 1/2$, if pairs of agents that reported $(\overline{v}, \overline{c})$ and $(\underline{v}, \underline{c})$ are present, an increase in the designer's payoff is achieved by rematching these pairs to create a $(\overline{v}, \underline{c})$ pair that generates a gain of 1 rather than $2\Delta_\alpha$.

Second, as the designer's problem is to determine which pairs should be cleared from the market in each period, when a pair that reported $(\overline{v}, \overline{c})$ or $(\underline{v}, \underline{c})$ is present, the designer has an incentive to wait (rather than clear the market) in the hope of eventually rematching pairs to create a $(\overline{v}, \underline{c})$ trade. In principle, this decision depends on the entire history of agent reports. However, as shown below, the state space can be simplified considerably. We make use of the fact that the relaxed optimization problem can be rewritten in terms of a Markov decision process. In Online Appendix B we provide a general description of the Markov decision process methodology developed in this paper.

We call a $(\overline{v}, \underline{c})$ pair *efficient*, the pairs $(\overline{v}, \overline{c})$ and $(\underline{v}, \underline{c})$ *suboptimal* and a $(\underline{v}, \overline{c})$ pair

[20] This also makes the incentive problem trivial: Buyers can simply be asked to pay $\overline{v}$ if they trade and sellers can be paid $\underline{c}$ if they trade.

inefficient. The underlying state at time t is identified as follows. We first determine the number of efficient pairs present and then determine the number of identical suboptimal $(\overline{v}, \overline{c})$ or $(\underline{v}, \underline{c})$ pairs present among the remaining set of agents. Our observations above imply that it cannot be optimal that non-identical suboptimal pairs, $(\overline{v}, \overline{c})$ and $(\underline{v}, \underline{c})$, are simultaneously present as these pairs can be split and rematched to form one efficient pair and one inefficient pair. Inefficient pairs can be ignored since these do not generate positive surplus and cannot be rematched to create efficient pairs. Thus, the *state space* of the designer's Markov decision process is two-dimensional and given by $\mathcal{X} := \{(x_E, x_S) : x_E, x_S \in \mathbb{Z}_{\geq 0}\}$, where x_E and x_S are the number of efficient pairs and suboptimal pairs present, respectively. Let $\boldsymbol{X}_t \in \mathcal{X}$ denote the state of the market after the arrival of period t agents.

We let $\mathcal{A}_{\boldsymbol{x}}$ denote the set of *actions* available to the designer in state $\boldsymbol{x}$, and $\mathcal{A} = \cup_{\boldsymbol{x} \in \mathcal{X}} \mathcal{A}_{\boldsymbol{x}}$. We have $\mathcal{A}_{\boldsymbol{x}} = \{(a_E, a_S) : a_E, a_S \in \mathbb{Z}_{\geq 0}, a_E \leq x_E, a_S \leq x_S\}$, where a_E and a_S denote the respective number of efficient pairs and suboptimal pairs being cleared from the market. Let $\boldsymbol{A}_t$ denote the action taken by the designer in period $t \in \mathbb{N}$, and denote by

$$P_{\boldsymbol{a}}(\boldsymbol{x}, \boldsymbol{y}) := \mathbb{P}(\boldsymbol{X}_{t+1} = \boldsymbol{y} \mid \boldsymbol{X}_t = \boldsymbol{x}, \boldsymbol{A}_t = \boldsymbol{a})$$

the *transition probability* that, if the designer takes the action $\boldsymbol{a}$ in state $\boldsymbol{x}$ in period t, the state in period $t+1$ will be $\boldsymbol{y}$. For any action $\boldsymbol{a} = (a_E, a_S)$, we have

$$P_{\boldsymbol{a}}(\boldsymbol{x}, (x_E - a_E + 1, x_S - a_S)) = p^2 \quad \text{and} \quad P_{\boldsymbol{a}}(\boldsymbol{x}, (x_E - a_E, x_S - a_S)) = (1-p)^2.$$

If $x_S = 0$ or $a_S = x_S$, a suboptimal pair arriving in period $t+1$ cannot be rematched. We have

$$P_{\boldsymbol{a}}(\boldsymbol{x}, (x_E - a_E, 1)) = 2p(1-p).$$

Otherwise, if an identical suboptimal pair arrives, it cannot be rematched and if a non-identical suboptimal pair arrives, the efficient agents in each pair can be rematched to form one efficient pair. Consequently, we have

$$P_{\boldsymbol{a}}(\boldsymbol{x}, (x_E - a_E, x_S - a_S + 1)) = P_{\boldsymbol{a}}(\boldsymbol{x}, (x_E - a_E + 1, x_S - a_S - 1)) = p(1-p).$$

We denote by

$$r(\boldsymbol{a}) = a_E + \Delta_\alpha a_S$$

the immediate *reward* when action $\boldsymbol{a} \in \mathcal{A}$ is implemented.

Given a Markov decision process $\langle \mathcal{X}, \mathcal{A}, P, r, \delta \rangle$, a *policy* $\pi : \mathcal{X} \to \mathcal{A}$ is such that $\pi(\boldsymbol{x}) \in \mathcal{A}_{\boldsymbol{x}}$ specifies the action taken by the designer in state $\boldsymbol{x}$. The *optimal policy* π^* of this Markov decision process maximizes the expected discounted reward earned by the designer, which by construction is given by (8). Thus, the designer's relaxed optimization problem reduces to determining the optimal policy π^* of $\langle \mathcal{X}, \mathcal{A}, P, r, \delta \rangle$. Since the state space $\mathcal{X}$ is countable, the feasible action sets $\mathcal{A}_{\boldsymbol{x}}$ are finite for all states $\boldsymbol{x}$ and the reward function is deterministic, a stationary deterministic optimal policy exists and is characterized by the appropriate Bellman equation (see, for example, Theorem 6.2.6 and Theorem 6.2.10 of Puterman (1994)).

3.2 Threshold policies and implementation

To determine the optimal policy we begin by defining a simple class of policies, which we call *threshold policies*. Threshold policies immediately clear efficient pairs from the market. Identical suboptimal pairs are stored up to a threshold $\tau \in \mathbb{N}$, and any additional suboptimal pairs are cleared immediately from the market.

Definition 1. *Given a threshold $\tau \in \mathbb{N}$, the associated* threshold policy π_τ *of the Markov decision process $\langle \mathcal{X}, \mathcal{A}, P, r, \delta \rangle$ is such that*

$$\pi_\tau(x_E, x_S) = (x_E, 0) \quad \text{if} \quad x_S \leq \tau \quad \text{and} \quad \pi_\tau(x_E, x_S) = (x_E, x_S - \tau) \quad \text{if} \quad x_S > \tau.$$

We now prove that the optimal market clearing policy is a threshold policy. This is intuitive, given that the designer essentially faces a binary choice in each period[21] and that the arrival process is stationary. Threshold policies are analogous to policies induced by a Gittins (1979) index, that apply to multi-armed bandit problems. Moreover, we show that it can be implemented using a P-IC and P-IR mechanism.

Theorem 1. *The optimal market clearing policy is a threshold policy that can be implemented with a P-IC and P-IR mechanism.*

[21] It is clearly optimal to immediately clear efficient pairs from the market so in each period the designer simply has to decide whether to clear or store after the arrival of an identical suboptimal pair.

We now briefly discuss uniqueness of the optimal allocation rule. Observe that the allocation rule that implements the optimal market clearing policy π^* is unique only up to the treatment of $(\underline{v}, \overline{c})$ pairs (that is, whether they are cleared or kept in the order book) and the identities of agents that are cleared from the market when more than one agent of a given type is present. Thus, in expressing the designer's optimization problem as a Markov decision process with a simple state space, we have shown that, given a market clearing policy, the designer's payoff does not vary with the treatment of individual agents. It immediately follows that the designer can implement the optimal market clearing policy using any queueing protocol over stored traders. For example, a first-come-first-served or a last-come-first-served queueing protocol could be used.[22] Note that since the queueing protocol serves as a tie-breaking rule, we can restrict attention to deterministic queueing protocols.

Corollary 1. *Given an optimal market clearing policy, social welfare gains and the profit of the designer are invariant to the treatment of inefficient pairs and the queueing protocol selected by the designer.*

Although Theorem 1 does not immediately allow us to identify the optimal market clearing policy, it is useful because it allows us to restrict attention to a small class of market clearing policies. This gives rise to a tractable dynamic programming approach, which we use to characterize the optimal threshold τ^*. In particular, each threshold policy π_τ induces a Markov chain $\{Y_t\}_{t\in\mathbb{N}}$ over $\{0, \ldots, \tau\}$, the number of identical suboptimal pairs stored in the order book. As is illustrated in Figure 2, $\{Y_t\}_{t\in\mathbb{N}}$ is a finite birth-and-death process. Computing the stationary distribution of this Markov chain is straightforward.

Proposition 2. *The stationary distribution $\boldsymbol{\kappa}$ of the Markov chain $\{Y_t\}_{t\in\mathbb{N}}$ under the threshold policy π_τ is given by*

$$\kappa_0 = \frac{1}{2\tau+1} \quad \text{and} \quad \kappa_i = \frac{2}{2\tau+1}, \; \forall i \in \{1, \ldots, \tau\}.$$

The designer's stationary expected per period payoff is given by

$$p^2 + \frac{2p(1-p)(\Delta_\alpha + \tau)}{2\tau+1}.$$

[22]We will later see that this invariance does not hold if we restrict the flexibility with which the designer sets transfers.

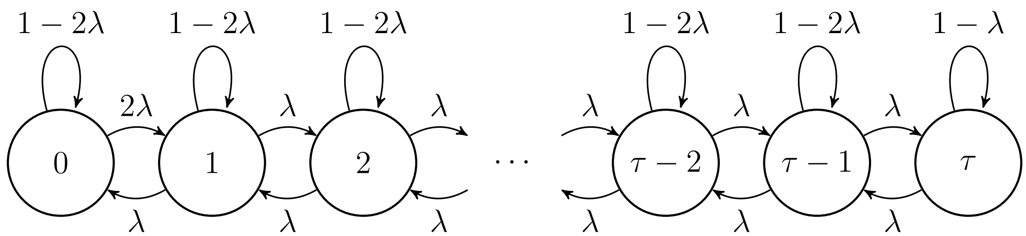

Figure 2: The Markov chain over the number of stored suboptimal pairs induced by the optimal policy π^*, where $\lambda = p(1-p)$.

The expression for the expected per period payoff has a simple and intuitive explanation. With probability p^2 an efficient pair arrives and trades, creating a welfare gain of 1. With probability $2p(1-p)$ a suboptimal pair arrives and there are several possibilities. With probability $(1/2)(1-\kappa_0) = \tau/(2\tau+1)$ there is a positive number of stored suboptimal pairs of the opposite kind. This arrival and the stored traders permit the creation of an efficient pair that trades and adds a welfare gain of 1. With probability $(1/2)\kappa_\tau = 1/(2\tau+1)$, τ suboptimal pairs of the same kind are stored, meaning that one suboptimal pair is cleared, generating a gain of Δ_α. In all other cases, the arriving suboptimal pair is stored and no immediate reward is earned by the designer.

Next, take any $y \in \{0, 1, \ldots, \tau\}$ and let $V_\tau^D(y)$ denote the expected present value of having y identical suboptimal pairs stored at the end of any period under the threshold policy with threshold $\tau \in \mathbb{N}$. Any such policy is, for $y \in \{1, \ldots, \tau - 1\}$, characterized by the Bellman equation

$$V_\tau^D(y) = \delta \left[p^2(1 + V_\tau^D(y)) + p(1-p)(1 + V_\tau^D(y-1) + V_\tau^D(y+1)) + (1-p)^2 V_\tau^D(y)\right], \quad (13)$$

with boundary conditions

$$V_\tau^D(0) = \delta \left[p^2(1 + V_\tau^D(0)) + 2p(1-p)V_\tau^D(1) + (1-p)^2 V_\tau^D(0)\right] \quad (14)$$

and

$$V_\tau^D(\tau) = \delta \left[p^2(1 + V_\tau^D(\tau)) + p(1-p)(1 + V_\tau^D(\tau - 1) + \Delta_\alpha + V_\tau^D(\tau)) + (1-p)^2 V_\tau^D(\tau)\right].$$

The optimal threshold τ^* can be determined using the stopping condition

$$V_{\tau^*}^D(\tau^*) > \Delta_\alpha + V_{\tau^*}^D(\tau^* - 1) \quad \text{and} \quad V_{\tau^*+1}^D(\tau^* + 1) \leq \Delta_\alpha + V_{\tau^*+1}^D(\tau^*). \quad (15)$$

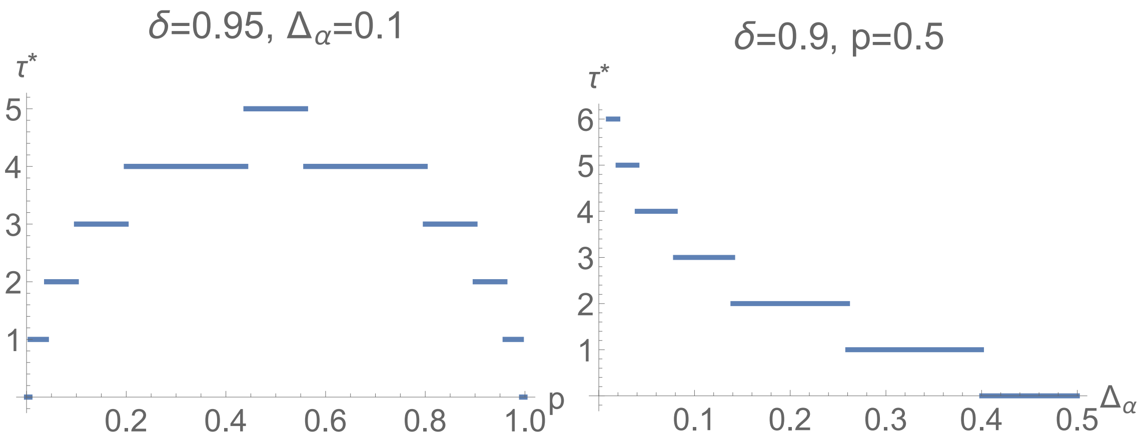

(a) Increasing in $p(1-p)$ (b) Decreasing in Δ_α

Figure 3: A numerical illustration of the comparative static results for τ^*.

To compute the optimal threshold, one can start with the threshold policy given by $\tau = 1$, check condition (15) and iterate. Algorithm C1 in Online Appendix C formalizes this procedure.

Proposition 3. *The optimal threshold τ^* is increasing in $p(1-p)$ and decreasing in Δ_0.*

Intuitively, τ^* increases as the cost of storing traders decreases and so is decreasing in Δ_α. Furthermore, the market maker stores suboptimal pairs in order to rematch them with identical suboptimal pairs in future periods. Thus, τ^* is increasing in the probability of such rematching, which, in a given period, is $p(1-p)$. Of course, τ^* is a straightforward measure of market thickness. These comparative statics are illustrated in Figure 3. Interestingly, Proposition 3 has the following corollary.

Corollary 2. *Market thickness, measured by τ^*, is increasing in α.*

Corollary 2 is reminiscent of Hotelling's (1931) finding that a monopolist extracts an exhaustible resource at a slower rate than a perfectly competitive industry.[23] As is the case

[23] Corollary 2 does not necessarily extend to finite horizon models with richer type spaces. For example, consider a two-period version of Myerson and Satterthwaite (1983) in which in every period a buyer-seller pair arrives, with a common discount factor applied to period two and with each agent drawing her type independently from a continuous distribution with compact support. Based on static mechanism design intuition, one might expect the market designer to increase profit by restricting trade in each period. However, this leads to a decrease in the probability that period one agents trade in period two, which reduces the benefit of waiting in period one. Thus, in some cases it is optimal for the market designer to increase period one trade to raise additional profit. See Loertscher et al. (2017)

in static environments, these distortions arise under the optimal mechanism as a means of reducing the informational rents of agents.[24]

Corollary 2 has important implications. In the perfectly patient limit (that is, as $\delta \to 1$), which, as noted, is equivalent to a static setup with a continuum of traders, a trade is executed if and only if it is efficient.[25] Consequently, the outcome, illustrated in Figure 1 in Section 2, is efficient and the average quantity traded per period (the Walrasian quantity) is p. Therefore, in static setups, suboptimal trades are indicative of inefficiency and possibly of rent extraction.[26]

However, the efficient outcome is different when $\delta < 1$. Under the efficient policy (that is, under π_{τ^*} for $\alpha = 0$), a suboptimal trade takes place in a given period if and only if a suboptimal pair arrives to a market in which τ^* identical suboptimal pairs are stored. Thus, $(\overline{v}, \overline{c})$ and $(\underline{v}, \underline{c})$ trades take place in each period with probability $p(1-p)/(2\tau^* + 1)$ as illustrated in Figure 4. Therefore, trades that are inefficient in a static setting are an integral part of efficiency in a dynamic setting. Moreover, keeping fixed the discount factor, by Corollary 2, the more such apparently inefficient suboptimal trades occur, the smaller is the market maker's rent extraction.

Another fundamental difference to static setups is that here efficiency is not a distribution-free concept because the optimal mechanism depends on p.

3.3 Implications for indirect taxation

The effect of different forms of indirect taxes on economic outcomes is another question of interest to economists, which has received renewed attention in the debates following the Global Financial Crisis about alternative forms of transaction taxes for financial markets. The question is also of relevance in policy debates pertaining to the remuneration scheme for artists whose songs are played by online streaming services such as Spotify, Pandora, or

[24] Inefficiently few matches also take place under profit maximization in the dynamic matching model of Fershtman and Pavan (2017).

[25] In the limit as $\delta \to 1$, there is no opportunity cost associated with storing suboptimal pairs and we must have $\tau^* \to \infty$. In the limit, all efficient agents are eventually cleared from the market as part of an efficient trade and inefficient agents trade with probability 0.

[26] For example, Yavaş (1996) investigates whether profit-seeking real estate brokers who earn a commission per trade have an incentive to maximize the number of trades, which is achieved by exclusively inducing suboptimal trades, rather than surplus.

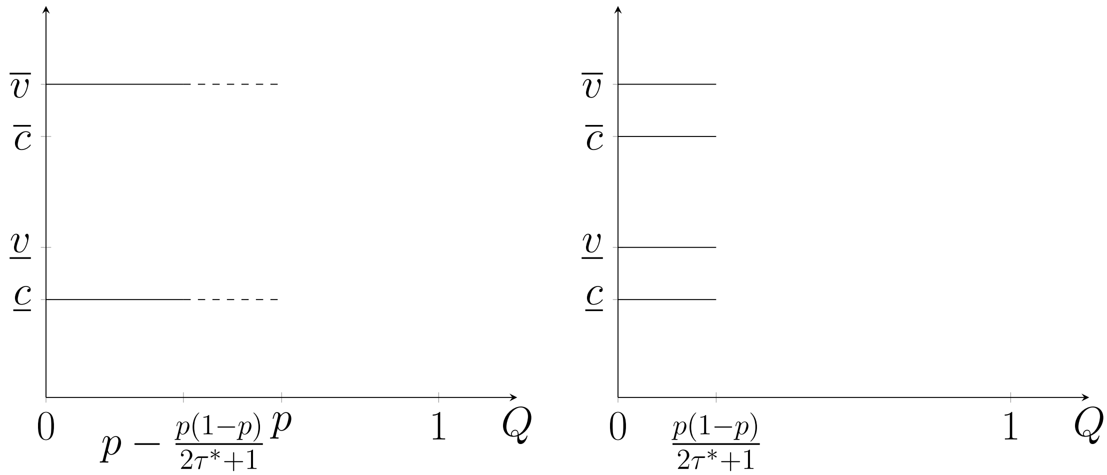

(a) Efficient trade

(b) Suboptimal trades

Figure 4: Efficiency for $\delta < 1$.

Apple, with some arguing that the platforms should be charged a fixed fee per song they play, which roughly corresponds to a specific tax, and others arguing that the platforms should be charged a percentage of their revenue, which can be interpreted as an ad valorem tax.

It is well known that, for perfectly competitive and thick markets, which in our setup correspond to the limit case as $\delta \to 1$, specific and ad valorem taxes are equivalent. In contrast, how these tax instruments compare in markets whose thickness is endogenously determined is an open question. To answer it, we now assume that the market maker is a profit maximizer and that authorities can observe and, under an ad valorem tax, tax the market maker's revenue.[27] This is analogous to the standard assumption in oligopoly models of indirect taxation that firms' profits can be observed and taxed.[28]

Under a specific tax $\sigma > 0$ per unit traded, the value of an efficient trade decreases from 1 to $1 - \sigma$ while the value of a suboptimal trade decreases from Δ_1 to $\Delta_1 - \sigma$. Given σ, the optimal policy of the designer is thus the same as for the our original problem with Δ_1 replaced by $\Delta(\sigma) = (\Delta_1 - \sigma)/(1 - \sigma)$. Observe that $\Delta' < 0$ and $\Delta(0) = \Delta_1$. Corollary 2, with Δ_α replaced by $\Delta(\sigma)$, thus implies that increasing σ will induce the market maker to

[27]We focus on profit-maximizing market makers in this subsection to simplify the analysis. Otherwise, we would have to derive the optimal policies and mechanisms anew and impose an assumption as to how much the market maker cares for tax revenue relative to social surplus and her own profit.

[28]This analysis extends directly to uniform and fixed frequency market clearing, which are introduced in Section 5 below.

increase the threshold τ^*. Thus, a specific tax distorts the relative value of suboptimal trades, inducing the market maker to create an excessively thick market and further reducing the welfare gains of buyers and sellers. When $\sigma > \Delta_1$, the market maker will become perfectly patient and never execute a suboptimal trade.

In contrast, an ad valorem tax levied as a percentage on the market maker's revenue will not affect the relative value of a suboptimal trade. Thus, the market clearing policy employed by the market maker will not change and an ad valorem tax can be levied without affecting social welfare gains. Consequently, we conclude that ad valorem taxes are superior to specific taxes in markets whose thickness is endogenously determined by a profit-maximizing market maker.[29]

3.4 General discrete type spaces

In Online Appendix B we show that the Markov decision process methodology is flexible and can be used to analyze a wide variety of extensions. In particular, we construct the Markov decision process for general discrete type spaces in 2.1 and prove appropriately generalized versions of Theorem 1, Proposition 3 and Corollary 2 in Theorem B1, Proposition B1 and Proposition B2 respectively. Note that our observations regarding indirect taxation also apply to general discrete type spaces.

Finally, in Online Appendix D, we consider several generalizations of the arrival process, including unpaired arrivals, continuous-time arrivals, group arrivals and multi-unit traders.

4 Posted prices

The optimal direct mechanism asks agents to report types and makes payments and allocations that depend, in general, on the reports of the contemporaneously arriving agents. In practice, often simpler, indirect mechanisms are used such as posted prices, which motivates us to examine this class of mechanisms. Doing so sheds new light on the possibility of efficient trade and allows us to derive equilibrium price distributions. It also enables us to be precise about the ways in which posted price mechanisms with a fixed spread are suboptimal

[29]Observe that the distorting effects of specific taxes vanish as δ approaches 1 because in the limit suboptimal trades vanish.

for a market maker who values profit. We begin with the following simple definition of a *posted price mechanism.*

Definition 2. *A* posted price mechanism *proceeds as follows. At the start of each period t, the designer posts a price p_B for buyers and a price p_S for sellers as a function of the state of the order book. The period t agents then arrive and all agents observe the order book and the posted prices before making a report of $\rho \in \{0, 1/2, 1\}$, where 1 indicates that the agent accepts the posted price, 0 indicates that the agent rejects the posted price and a report of 1/2 expresses that the agent is indifferent between accepting and rejecting the posted price. The designer then clears the market at the posted prices on the basis of these reports. In the event that there are ties, a queueing protocol specifies how these are broken.*

According to our definition, a posted price mechanism is thus characterized by the pricing rules p_B and p_S and the queueing protocol. We say that buyers (sellers) report truthfully under a posted price mechanism if they report $\rho = 1$ whenever $v > p_B$ ($c < p_S$), $\rho = 1/2$ whenever $v = p_B$ ($c = p_S$) and $\rho = 0$ whenever $v < p_B$ ($c > p_S$). To analyze the incentive properties of posted price mechanisms, we make the simplifying assumption that the designer removes from the market any agent that makes reports that are not consistent with truthful reporting for any type. Observe that under a truthful equilibrium, the designer can infer the stored traders' types.

4.1 An efficient budget balanced posted price mechanism

Consider now the following posted price mechanism, which we refer to as the *balanced budget posted price mechanism* associated with a threshold $\tau \geq 1$, defined as follows: If the number of suboptimal pairs stored in the order book is $y < \tau$, the designer posts prices $p_B = p_S = 1/2$. If the number of suboptimal pairs stored in the order book is $y = \tau$, the designer posts prices $p_B = p_S = \Delta_0$, provided the stored pairs are of the types $(\underline{v}, \underline{c})$, and the prices $p_B = p_S = 1 - \Delta_0$ if the stored pairs are of the types $(\overline{v}, \underline{c})$. A last-come-first-served queueing protocol is used to determine the order in which agents are cleared from the market and to break ties if necessary.

By construction, the balanced budget posted price mechanism does not run a deficit. In equilibrium, the mechanism immediately executes efficient trades and does not execute any suboptimal trades when less than τ identical suboptimal pairs are stored. Once τ pairs of type $(\overline{v}, \overline{c})$ are stored, the designer posts period t prices of $p_B = p_S = 1 - \Delta_0$, so that any efficient or additional suboptimal trades created in period t are executed. Similarly, once τ pairs of type $(\underline{v}, \underline{c})$ are stored, the designer posts period t prices of $p_B = p_S = \Delta_0$, so that any efficient or additional suboptimal trades created in period t are executed. Therefore, under truthful reporting the balanced budget posted price mechanism implements a threshold policy with threshold $\tau \geq 1$. When δ is equal to or close to 0, such a mechanism is not efficient because for δ sufficiently small, the efficient policy executes all trades that generate positive surplus. However, the downside to the efficient mechanism when δ is small and no trades are stored is that, depending on the parametrization, it may run a deficit.[30] Interestingly, it turns out that there is a tight connection between the budget balanced posted price mechanism and the efficient allocation rule as stated in the following proposition.

Proposition 4. *The following statements are equivalent: (i) The efficient allocation rule can be implemented using a P-IC and P-IR budget balanced posted price mechanism, and (ii)* $\tau^* > 0$ *for* $\alpha = 0$.

We now briefly develop the intuition behind this result. As discussed, the budget balanced posted price mechanism implements the efficient allocation rule if $\tau^* > 0$, provided agents report truthfully. As such, we only need to check the agents' incentive constraints. For agents of type $\underline{v}$ and $\overline{c}$, there is no incentive to misreport because these agents always receive a payoff of zero when reporting truthfully (regardless of the history and the types of contemporary agents) and cannot receive a positive expected discounted payoff by misreporting. Agents of type $\overline{v}$ will clearly report truthfully upon observing a price of Δ_0. If a price of $1/2$ or $1 - \Delta_0$ is observed, misreporting guarantees that the buyer will eventually leave the market without trading (either immediately or later due to the last-come-first-served queueing protocol) regardless of the history and the types of contemporary agents. Thus,

[30]This static mechanism is readily derived. The individual rationality constraints for the inefficient traders are made binding by making interim expected payments for the buyer of type $\underline{v}$ equal to $p\underline{v}$ and the interim expected payment to the seller of type $\overline{c}$ equal to $p\overline{c}$. Bayesian incentive compatibility for the efficient types then means that the buyer of type $\overline{v}$ pays no more than $(1-p)\overline{v} + p\underline{v}$ and the seller of type $\underline{c}$ is paid not less than $p\overline{c} + (1-p)\underline{c}$. Substituting $\overline{v} = 1$ and $\underline{c}$, the maximized expected revenue of the market maker, subject to efficiency, incentive compatibility and individual rationality constraints, is thus $p(2\Delta_0 - p)$, which is negative for $p > 2\Delta_0$.

the incentive constraints are satisfied for buyers of type $\overline{v}$. A similar argument applies to sellers of type $\underline{c}$.

Recall that in our discussion of the optimal mechanism we noted that, due to a version of the revenue equivalence theorem, any queueing protocol can be used to break ties. However, under a posted price mechanism, the choice of queueing protocol matters because we have less flexibility in determining transfers.

We immediately have the following corollary to Proposition 4:

Corollary 3. *The efficient mechanism does not run a deficit if* $\tau^* > 0$.

As noted in Footnote 30, any efficient, incentive compatible and individually rational mechanism runs a deficit when $\delta = 0$ if $\Delta_0 < p/2$. Proposition 4 thus sheds new light on the impossibility of efficient trade along the lines of Myerson and Satterthwaite (1983) for dynamic environments.[31] Indeed, dynamics and optimally trading off gains from market thickness against costs of delay offer a way of overcoming the impossibility of efficient trade. On the surface, this is related to the strand of literature in the tradition of Gresik and Satterthwaite (1989) that investigates how quickly inefficiency disappears as the number of buyers and sellers increases in markets that are constrained not to run a deficit. However, our result does not merely or primarily rely on a large markets argument. In our setting, the deficit vanishes as soon as $\tau^* > 0$, which occurs for $\delta = 0.36$ for the parameters used in Figure 3.[32] Remarkably, the price posting implementation permits efficiency, not only without running a deficit in expectation, but in fact with a balanced budget at all times.[33]

[31] The impossibility result of Myerson and Satterthwaite (1983) does not hold as generally for binary type distributions as it does for continuous distributions. See Matsuo (1989) for a treatment of the bilateral problem of Myerson and Satterthwaite (1983) with binary types and Kos and Manea (2009) for a version with general discrete types.

[32] Given some value of δ, the expected (or discounted) number of pairs present in our setting would be $1/(1-\delta)$. For δ in the order of 0.36 and $\Delta = 0.1$ and $p = 0.5$, the dynamic efficient mechanism does not run a deficit (see Corollary 3 and Figure 3). Because $1/(1-0.36) \approx 1.5$, the expected number of pairs present is less than 2. The parametrization $p > 2\Delta_0$, which is sufficient to have a deficit in static, ex post efficient bilateral trade (see footnote 30), is also sufficient for a deficit with $N = 2$ pairs present.

[33] As noted in the introduction (see in particular footnote 3), price posting and efficiency also go hand in hand in static bilateral trade problems. With independent, continuous distributions with overlapping supports, Myerson and Satterthwaite (1983) prove the impossibility of ex post efficient trade subject to incentive compatibility and individual rationality. Consequently, a posted price mechanism, while balancing the budget, constrains social surplus. With non-overlapping supports, ex post efficiency is possible and can be implemented with a posted price (for example, by setting the price equal to the mid-point between the lower (upper) bound of support of the buyer's (seller's) distribution).

To see intuitively why implementation of the efficient allocation is possible whenever $\tau^* \geq 1$, notice that an arriving inefficient agent trades upon arrival with non-zero probability only if they are from a particular side of the market (inefficient agents trade only if they are on the side of the market for which no efficient agents have accumulated). Thus, in each period the designer knows a priori whether the buyer's price must be sufficiently low so that a $\underline{v}$ type may trade if necessary or the seller's price must be sufficiently high so that a $\overline{c}$ type may trade if necessary. This is also precisely why implementation using a posted price mechanism with $p_B = p_S$ is possible.

Proposition 4 uncovers a dynamic connection between two fundamental theorems in economics: The *Coase Theorem* and the *Myerson-Satterthwaite Theorem.* The former states that, absent transaction costs, efficient trade of resource poses no problem and the latter that, in a static bilateral trade setting, private information poses an insurmountable obstacle for efficient trade. Proposition 4 shows that efficient, budget-balanced trade is possible if storing enough traders is efficient. Put differently, in our setup the Coase Theorem applies if there is an appropriately designed dynamic market mechanism and agents are not too impatient. Importantly, the main conclusion relies on dynamic regularity (as defined in (5)) and does not depend on the type space being binary, as shown in the following:

Proposition 5. *Take any dynamically regular discrete type spaces such that, for $\delta = 0$, any ex post efficient, incentive compatible and individually rational mechanism runs a deficit. Then there exists a sufficiently large value of δ that is less than 1 such that, under the efficient mechanism that satisfies P-IC and P-IR, the designer's expected discounted profit is positive.*

The result that the efficient mechanism can be implemented via price posting for $\delta < 1$ does not easily generalize to richer type spaces. The analysis of Gresik and Satterthwaite (1989) shows that, for continuous type spaces, the efficiency loss associated with the no-deficit constraint only vanishes in the limit as $\delta \to 1$.

4.2 Equilibrium price distribution

The implementation of the efficient policy, provided $\tau^* > 0$, via the simple posted price mechanism also enables us to characterize the stationary price distribution and to provide a

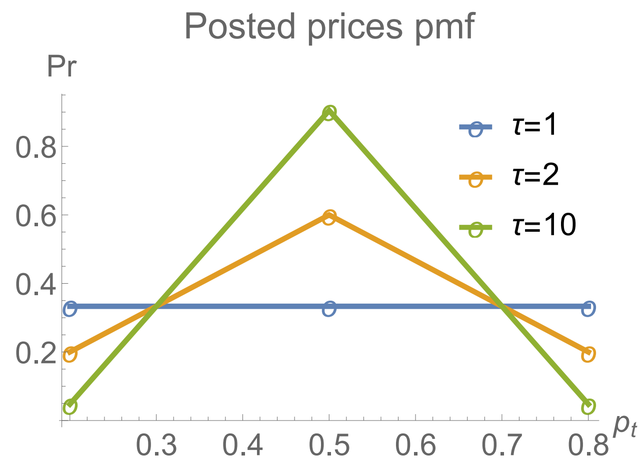

Figure 5: The probability mass function of the posted price under the budget balanced posted price mechanism for $\Delta_0 = 0.2$ implying $P_t \in \{0.2, 0.5, 0.8\}$.

measure of *market thickness*, which we take to be an individual trader's likely price impact. We begin with the characterization of the steady state distribution. Let P_t denote the price posted in period t. Under the stationary distribution, we have

$$\mathbb{P}(P_t = \underline{v}) = \mathbb{P}(P_t = \overline{c}) = \frac{1}{2\tau^* + 1} \quad \text{and} \quad \mathbb{P}(P_t = 1/2) = \frac{2\tau^* - 1}{2\tau^* + 1}, \tag{16}$$

where the equalities in (16) follow from the stationary distribution given in Proposition 2. An illustration of this distribution is given in Figure 5. It follows that the stationary variance of the posted prices is

$$\text{Var}(P_t) = \frac{2}{2\tau^* + 1} \left(\frac{1}{2} - \Delta_0 \right)^2 .$$

Based on these formulas, we now have the following corollary to Proposition 4:

Corollary 4. $\mathbb{P}(P_t = 1/2)$ *increases in* δ *and* $p(1-p)$ *and decreases in* Δ_0*.* $\text{Var}(P_t)$ *decreases in* δ *and* $p(1-p)$*. Moreover,*

$$\lim_{\delta \to 1} \mathbb{P}(P_t = 1/2) = 1 \quad \text{and} \quad \lim_{\delta \to 1} \text{Var}(P_t) = 0.$$

The first part of the corollary says that the distribution shifts more weight to the static Walrasian price as the discount factor increases and the probability of a contemporaneous mismatch (that is, $p(1-p)$) increases. It decreases as the value of a suboptimal trade increases because the optimal threshold τ^* decreases in this value. Likewise, the price variance decreases in the discount factor and the probability of a mismatch. However, the effect of the value of a suboptimal trade on the price variance cannot be signed in general because, on the one hand, such increases shift probability mass to the extremes, thereby increasing the variance when all else is equal, while on the other hand they narrow the gap between the lowest price Δ_0 and the highest price $1 - \Delta_0$ in the support.

The limit results in the second part of Corollary 4 state that the equilibrium price distribution converges to a degenerate distribution that has probability 1 on the static Walrasian price of 1/2. This result resonates with classic convergence results in the literature on the microfoundation of competitive equilibrium such as Satterthwaite and Shneyerov (2007) or Lauermann (2013), which provide sufficient conditions for equilibrium in dynamic search and matching settings to converge to the (static) Walrasian equilibrium as search frictions (often also parametrized by a discount factor) vanish. However, there is a subtle but important difference: In the aforementioned papers, the equilibrium allocation is inefficient for $\delta < 1$ whereas in our setting, equilibrium behavior under the posted price mechanism is, by construction of the mechanism, efficient for any δ, provided only it is large enough so that $\tau^* > 0$.

We now turn to the determination of an individual agent's likely price impact, which can be interpreted as a measure of market thickness. In so doing, we stipulate that an agent arrives to an order book that is characterized by the stationary distribution in period t and ask what is the probability that this agent's truthful reporting changes the price from the static Walrasian price of 1/2 to one of the two extremes ($\bar{c}$ for a buyer and $\underline{v}$ for a seller). Notice that an agent only has a price impact, given $p_{t-1} = 1/2$, if the number of identical suboptimal pairs in the order book is at the threshold value τ^* *and* if she is part of another, identical suboptimal pair. Therefore, defined in this way, an agent's likely price impact is

$$p_{im} := \mathbb{P}(P_t = \underline{v} | p_{t-1} = 1/2) = \mathbb{P}(P_t = \bar{c} | p_{t-1} = 1/2) = \frac{p(1-p)}{2\tau^* + 1}.$$

Proposition 4 implies that p_{im} decreases in δ and increases in Δ_0. That is, the greater is the discount factor (the smaller is the value of a suboptimal trade), the smaller is an individual agent's likely price impact (and the thicker is the market, measured in this way). Whether p_{im} increases or decreases in the probability $p(1-p)$ of a contemporaneously arriving suboptimal pair cannot be determined in general because of the two opposing effects: the threshold τ^* increases in $p(1-p)$, which all else decreases p_{im}, but $p(1-p)$ directly increases p_{im} because a suboptimal pair is required to move the price away from 1/2 in the first place.

4.3 The profit-maximizing efficient posted price mechanism

Extending our previous analysis, we now consider posted price mechanisms under which the designer charges traders a bid-ask spread and compute the posted price mechanism that implements the efficient allocation, whilst maximizing profit for the market maker.

We define the *bid-ask spread posted price mechanism* associated with a threshold $\tau \geq 1$ as follows: If the number of suboptimal pairs stored in the order book is $y < \tau$, the designer posts prices $p_B = 1$ and $p_S = 0$. If the number of suboptimal pairs stored in the order book is $y = \tau$, the designer posts prices $p_B = \Delta_0$ and $p_S = 0$, provided the stored pairs are of the types $(\underline{v}, \underline{c})$, and the designer posts the prices $p_B = 1$ and $p_S = 1 - \Delta_0$ otherwise. A last-come-first-served queueing protocol is used to determine the order in which agents are cleared from the market and to break ties if necessary.

Proposition 6. *The bid-ask spread posted price mechanism with $\tau = \tau^*$ provides a P-IC and P-IR implementation of the efficient allocation and, within the class of posted price mechanisms identified in Definition 2, maximizes the market maker's profit.*

The intuition behind this result is as follows. If the designer increases any of the posted prices in any period, the mechanism will fail to implement the efficient allocation. So we have found the desired mechanism provided the appropriate incentive constraints hold. Checking the P-IR constraints and the P-IC constraints for worst-off types is completely routine, while the P-IC constraints hold for the efficient types by virtue of the last-come-first-served queueing protocol, which ensures that efficient types that misreport eventually leave the market without trading.[34] The bid-ask spread posted price mechanism, although optimal among posted price mechanisms, does not coincide with the profit-maximizing efficient mechanism

[34] Note that with a different queueing protocol (such as rationing uniformly at random), the designer would not necessarily be able to offer $p_B = 1$ and $p_S = 0$ whenever the order book is below capacity. If there is a non-zero probability that a buyer of type $\overline{v}$ who rejected $p_B = 1$ is eventually able to trade at $p_B = \Delta_0$, the market maker would need to offer a spread with $p_B < 1$ so that the incentive compatibility constraint for efficient buyers is not violated. Similar logic applies to sellers.

because of foregone profit on efficient trades executed with a buyer's price of $p_B = 1 - \Delta_0$ or a seller's price of $p_S = \Delta_0$.

Notice also that the profit maximizing, efficient bid-ask spread posted price mechanism involves a spread that varies with the order book. In particular, the spread is 1 when $y < \tau^*$ and Δ_0 otherwise. Thus, Proposition 6 implies any posted price with a fixed spread, or equivalently, with a fixed fee per unit traded, is not a profit-maximizing, efficient price posting mechanism.

5 Approximately optimal mechanisms

Beyond incentive compatibility, individual rationality and feasibility, in reality, additional constraints are often imposed on market makers. For example, in foreign exchange spot markets such as Thomson Reuters, ParFX, and EBS, clearing is *uniform* in that all compatible orders are cleared at once when clearing occurs while the time intervals that elapse between clearings depend on the orders received. In other trading venues, such as the Global Dairy Trade (GDT), market clearing is both uniform and occurs at a fixed frequency (fortnightly in the case of GDT), which we refer to as *fixed frequency* market clearing.

Definition 3. *Under* uniform market clearing *the entire market is cleared at the time of clearing.* Fixed frequency market clearing *requires that, in addition to market clearing being uniform, the market is cleared at fixed intervals.*

In this section we study the optimal mechanisms under these additional and natural restrictions. In particular, we show that these mechanisms are approximately optimal when the discount factor is large enough. The intuitive reason is that, as everyone becomes patient enough, the designer executes only high-margin trades, which is possible to approximate even under these natural restrictions. Furthermore, provided the agents are sufficiently patient, a profit-maximizing designer generates greater welfare than a more constrained, welfare-targeting designer. For ease of exposition we again consider the symmetric binary type setup in the section. However, in Online Appendix B.3 we show that the main results of this section (Theorems 2 and 3) generalize to general discrete type spaces (see Theorems B3 and B4 in the Online Appendix). Finally, we also defer the derivation and analysis of the class of optimal mechanisms under uniform and fixed frequency market clearing for binary type spaces to Online Appendix B.

Let $W^{D,\alpha}$ denote the expected discounted welfare gains (starting from an empty market at $t = 0$) under *discriminatory* market clearing, which corresponds to the form of market clearing we have studied thus far, with a designer who maximizes $(1-\alpha)W + \alpha R$ with $\alpha \in [0,1]$. Similarly, we use the notation $W^{U,\alpha}$, $W^{F,\alpha}$ and $W^{0,\alpha}$ for uniform market clearing, fixed frequency market clearing and instantaneous trade, respectively. Total expected discounted welfare gains from the periodic ex post efficient market mechanism, which for brevity we simply call *instantaneous trade*, are[35]

$$W^{0,0}(\delta) = \frac{1}{1-\delta}\left(p^2 + 2p(1-p)\Delta_0\right). \tag{17}$$

5.1 Asymptotic gains from sophistication

We first compare the benefits from increasing sophistication for a given α as $\delta \to 1$. Since welfare under each form of market clearing diverges in this limit, we consider the relative gains from increased sophistication that are defined as

$$G_\alpha^{D,U}(\delta) := \frac{W^{D,\alpha}(\delta) - W^{U,\alpha}(\delta)}{W^{D,\alpha}(\delta)} \quad \text{and} \quad G_\alpha^{U,F}(\delta) := \frac{W^{U,\alpha}(\delta) - W^{F,\alpha}(\delta)}{W^{U,\alpha}(\delta)}$$

and

$$G_\alpha^{F,0}(\delta) := \frac{W^{F,\alpha}(\delta) - W^{0,\alpha}(\delta)}{W^{F,\alpha}(\delta)}.$$

Theorem 2. *In the $\delta \to 1$ limit, the relative gains from sophistication are given by*

$$\lim_{\delta\to 1} G_\alpha^{D,U}(\delta) = \lim_{\delta\to 1} G_\alpha^{U,F}(\delta) = 0 < (1-p)(1-2\Delta_\alpha) = \lim_{\delta\to 1} G_\alpha^{F,0}(\delta).$$

Theorem 2 is illustrated numerically in Figure 6. The theorem says that the relative gains from additional sophistication vanish while the relative gains from any degree of sophistication relative to instantaneous trade remain strictly positive as δ approaches 1. Intuitively,

[35] It is interesting to note that under continuous-time double auction mechanisms feasible trades are also executed immediately. Thus, the outcome of instantaneous trade is the same as the outcome that would result under a continuous-time double auction with truthful bidding. Continuous-time double auctions are not incentive compatible as the bid of a given trader affects both the probability of trade and, in the event that trade occurs, the market price. Under strategic bidding one would expect efficient types to bid shade in order to avoid trading with an inefficient type so that they receive a higher expected payoff. Although the equilibrium behavior of a continuous-time double-auctions is difficult to characterize (see for example, Satterthwaite and Williams (2002)), the outcome under the first-best mechanism provides an efficiency benchmark for evaluating the outcome of a continuous-time double auction.

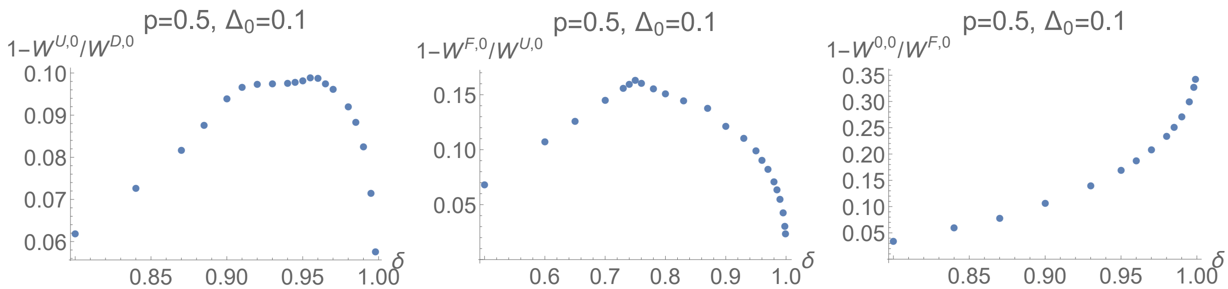

(a) Discrim. vs. uniform (b) Uniform vs. fixed freq. (c) Fixed freq. vs. decent.

Figure 6: The relative gains of additional sophistication.

because instantaneous trade corresponds to fixed frequency market clearing with the frequency given by the period and because all other forms of dynamic market clearing which we consider impose fewer restrictions than fixed frequency market clearing, it also follows that the welfare-maximizing mechanism under uniform and fixed frequency market clearing weakly outperforms instantaneous trade in terms of social welfare gains. The outcome under instantaneous trade coincides with the efficient outcome when δ is so close to 0 that storing is not efficient.

Much harder is the comparison of social welfare gains under the optimal *profit-maximizing mechanism* to those under instantaneous trade because a simple argument based on less constrained optimization cannot be used: As one goes from periodic ex post efficient trade to profit-maximizing discriminatory market clearing, one not only eliminates constraints but also alters the objective that is maximized. We provide this comparison in the following subsection.

5.2 Welfare under sophisticated, profit-targeting exchanges

Creating larger markets and employing increasingly sophisticated mechanisms may involve costs such as advertising and promotion, physical infrastructure investments, and labor. In addition, large exchanges are often operated by profit-seeking companies whereas small exchanges through which instantaneous trade occurs can plausibly be thought of as extracting little to no rents. This renders the question relevant whether a larger, profit-maximizing exchange creates greater welfare gains than instantaneous (periodic ex post efficient) trade. In the main result of this section, we show that the answer is affirmative, provided the discount factor is large enough. Further, a pertinent issue in the design of two-sided markets is the need to "bring both sides of the market on board" (see, for example, Caillaud and Jullien, 2003; Rochet and Tirole, 2006). While a full analysis of this question requires a different model and is thus beyond the scope of this paper, the following result sheds new light on this question.

Theorem 3. *For all $k \in \{D, U, F\}$ there exists $\delta_k \in [0, 1)$ such that $W^{k,1}(\delta_k) = W^{0,0}(\delta_k)$ and $W^{k,1}(\delta_k) > W^{0,0}(\delta_k)$ for all $\delta > \delta_k$.*

Theorem 3 says that for sufficiently large discount factors, a profit-maximizing platform generates greater welfare gains than instantaneous, periodic ex post efficient trade. This is so because efficient types trade with relatively high probability under the profit-maximizing platform, which is efficient for a sufficiently large discount factor. Therefore, if a profit-oriented centralized platform needs to attract buyers and sellers from a welfare-maximizing platform which induces trade instantaneously, by Theorem 3 the profit-oriented platform can do so by offering a sufficiently high share of the trade surplus to efficient types while extracting all surplus from the inefficient types. By getting the key players on board (in our setting, these are the buyers of type $\overline{v}$ and the sellers with cost $\underline{c}$) the others will have no choice but to follow suit.

One natural interpretation is that the profit-maximizing market maker is an Internet giant while the periodic ex post efficient exchange can be thought of as a brick-and-mortar retailer (or a mum-and-dad shop). Although the theorem does, of course, not imply that profit-maximizing Internet giants are necessarily better for social welfare than more traditional shops, it does provide a formalization of the notion of overwhelming returns to scale due to the gains from market thickness.

6 Conclusions

Economic agents interact in an inherently dynamic world. Agents without a trading partner today may find one in the future, and agents with a possible trading partner today may find better trading opportunities further down the track. While a large literature on the micro-foundations of Walrasian equilibrium has studied equilibrium behavior as (search) frictions (often captured by a discount factor) vanish, we address the converse question of what is the best a market maker can do for a given discount factor. To be specific, we derive the optimal market mechanism for an environment with stochastically arriving traders who are privately informed about their values and costs. This mechanism balances the gains from increased market thickness against the opportunity cost of delay from waiting to clear the market.

We show that, with binary types, efficient, incentive compatible and individually rational trade is possible with an ex post budget balanced mechanism – posted prices – if it is optimal to store at least one trade. This result has a Coasian flavor because it means that initial misallocations can be resolved efficiently if agents are not too impatient. At the same time, it also provides a rationale for market design because instantaneous trade (that is, a periodic ex post efficient mechanism that never stores traders) is not efficient in our dynamic setting under these conditions.

The distribution of posted prices that implement the efficient allocation rule is uniform when it is optimal to store one trade and converges to a degenerate distribution with all mass on a single Walrasian price for a static model with a continuum of traders as the discount factor approaches one. While these results are reminiscent of findings in the literature on the microfoundation of Walrasian equilibrium, there is an important difference: The distribution of posted prices we derive implements the efficient allocation rule for any discount factor, provided only it is optimal to store at least one trade.

We also derive the mechanism that maximizes the market maker's expected profit and show, among other things, that the social welfare gains associated with this mechanism exceed social welfare gains of a periodic ex post efficient mechanism that never stores trades if the discount factor is sufficiently large. While most of our analysis allows the market maker the clear the market discriminatorily, we extend the analysis for uniform and fixed frequency market clearing and show that, as the discount factor approaches one, the relative social welfare gains under all of these mechanisms converge.

Our paper opens a number of avenues for future research. For example, introducing product differentiation – such as spatial differentiation – would allow one to analyze dynamic allocation problems such as those that ride-sharing service providers face. Another natural and promising extension would be to endow agents with some units of the good and have them decide endogenously whether they want to buy or sell, which would permit a dynamic analysis of asset markets in which agents choose their trading positions – buy, sell, hold – endogenously.

Furthermore, a large number of papers in the finance literature have shown that large institutional traders optimally reduce their price impact by breaking up their trades when a fixed, suboptimal mechanism is used to clear the market.[36] In light of this, an interesting extension of our model would be to accommodate large traders and to analyze the impact of traders' size on the optimal market clearing mechanism. Given the prominent role interdependent values play in models of financial markets, it would also be valuable to allow for some form of interdependence in our setting. However, because it is not quite clear what the suitable set of assumptions would be when agents with persistent types arrive over time, this problem is best left for future research. More fundamentally, our paper offers the possibility of analyzing market thickness, price distributions, and market impact under the hypothesis that the market operates under an optimal mechanism for a variety of environments.

References

Akbarpour, M., S. Li, and S. O. Gharan (2017): "Thickness and Information in Dynamic Matching Markets," Working paper.

Anderson, R., I. Ashlagi, D. Gamarnik, and Y. Kanoria (2017): "Efficient Dynamic Barter Exchange," *Operations Research*, 65, 1446–1459.

Athey, S. and D. A. Miller (2007): "Efficiency in repeated trade with hidden valuations," *Theoretical Economics*, 2, 299–354.

Athey, S. and I. Segal (2013): "An Efficient Dynamic Mechanism," *Econometrica*, 81, 2463–2485.

Baccara, M., S. Lee, and L. Yariv (2016): "Optimal Dynamic Matching," Working paper.

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

[36]See, among others, Vayanos (1999), Rostek and Weretka (2015) and Du and Zhu (2017).

Barlow, R. E. and F. Proschan (1975): *Statistical Theory of Reliability and Life Testing*, New York: Holt, Rinehart and Winston.

Bergemann, D. and J. Välimäki (2010): "The Dynamic Pivot Mechanism," *Econometrica*, 78, 771–789.

Board, S. and A. Skrzypacz (2016): "Revenue Management with Forward-Looking Buyers," *Journal of Political Economy*, 124, 1046–1087.

Börgers, T. (2015): *An Introduction to the Theory of Mechanism Design*, New York: Oxford University Press.

Borovkov, K. (2014): *Elements of Stochastic Modelling*, World Scientific Publishing Company, 2 ed.

Budish, E., P. Cramton, and J. Shim (2015): "The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response," *Quarterly Journal of Economics*, 130, 1547–1621.

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

Caillaud, B. and B. Jullien (2003): "Chicken and Egg: Competition Among Intermediation Service Providers," *RAND Journal of Economics*, 34, 309–328.

Che, Y.-K. (2006): "Beyond the Coasian Irrelevance: Asymmetric Information," Unpublished lecture notes.

Coase, R. H. (1960): "The Problem of Social Cost," *Journal of Law and Economics*, 3, 1–44.

Cramton, P., R. Gibbons, and P. Klemperer (1987): "Dissolving a Partnership Efficiently," *Econometrica*, 55, 615–632.

Crémer, J. and R. P. McLean (1985): "Optimal Selling Strategies Under Uncertainty for a Discriminating Monopolist when Demands are Interdependent," *Econometrica*, 53, 345–362.

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

Delacrétaz, D., S. Loertscher, L. Marx, and T. Wilkening (2018): "Two-Sided Allocation Problems, Matching with Transfers, and the Impossibility of Ex Post Efficiency," *Journal of Economic Theory*, forthcoming.

Du, S. and H. Zhu (2017): "What is the Optimal Trading Frequency in Financial Markets?" *Review of Economic Studies*, 84, 1606–1651.

Elkind, E. (2007): "Designing and Learning Optimal Finite Support Auctions," in *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, Society for Industrial and Applied Mathematics, 736–745.

Fershtman, D. and A. Pavan (2017): "Matching Auctions," Working paper.

Garrett, D. F. (2016): "Ready to trade? On budget-balanced efficient trade with uncertain arrival," Working paper.

Gershkov, A. and B. Moldovanu (2010): "Efficient Sequential Assignment with Incomplete Information," *Games and Economic Behavior*, 68, 144–154.

Gittins, J. C. (1979): "Bandit Processes and Dynamic Allocation Indices," *Journal of the Royal Statistical Society, Series B*, 41, 148–177.

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

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

Hendricks, K. and A. Sorensen (2018): "Dynamics and Efficiency in Decentralized Online Auction Markets," Working paper.

Hotelling, H. (1931): "The Economics of Exhaustible Resources," *Journal of Political Economy*, 39, 137–175.

Hurwicz, L. (1972): "On Informationally Decentralized Systems," in *Decision and Organization*, ed. by C. B. McGuire and R. Radner, Amsterdam: North-Holland, 297–336.

Kelly, F. and E. Yudovina (2016): "A Markov Model of a Limit Order Book: Thresholds, Recurrence, and Trading Strategies," Working paper.

Kos, N. and M. Manea (2009): "Efficient Trade Mechanism with Discrete Values," Working paper.

Kosmopoulou, G. and S. R. Williams (1998): "The Robustness of the Independent Private Value Model in Bayesian Mechanism Design," *Economic Theory*, 12, 393–421.

Larsen, B. J. (2018): "The Efficiency of Real-World Bargaining: Evidence from Wholesale Used-Auto Auctions," Working paper.

Latouche, G. and V. Ramaswami (1999): *Introduction to Matrix Analytic Methods in Stochastic Modeling*, Philadelphia, Pennsylvania: Society for Industrial and Applied Mathematics.

Lauermann, S. (2013): "Dynamic Matching and Bargaining Games: A General Approach," *American Economic Review*, 103, 663–689.

Lerner, A. P. (1944): *The Economics of Control*, New York: Macmillan Company.

Lewis, M. (2014): "The Woof Hunters of Wall Street," *The New York Times Magazine*.

Loertscher, S. and L. Marx (2017): "Optimal Clock Auctions," Working paper.

Loertscher, S., E. V. Muir, and P. G. Taylor (2017): "Dynamic Market Making," Working paper.

Matsuo, T. (1989): "On Incentive Compatible, Individually Rational, and Ex Post Efficient Mechanisms for Bilateral Trading," *Journal of Economic Theory*, 49, 189–194.

McCrank, J. (2017a): "Chicago Stock Exchange plans new speed bump to curb fast traders," *Reuters*.

——— (2017b): "NYSE wins regulatory approval for 'speed bump' exchange," *Reuters*.

Melton, H. (2017): "Market Mechanism Refinmenet on a Continuous Limit Order Book Venue," Working paper.

Mendelson, H. (1982): "Market Behaviour in a Clearing House," *Econometrica*, 50, 1505–1524.

Mezzetti, C. (2004): "Mechanism Design with Interdependent Valuations: Efficiency," *Econometrica*, 72, 1617–1626.

Mierendorff, K. (2013): "The Dynamic Vickrey Auction," *Games and Economic Behavior*, 82, 192–204.

——— (2016): "Optimal Dynamic Mechanism Design with Deadlines," *Journal of Economic Theory*, 161, 190–222.

Milgrom, P. (2017): *Discovering Prices*, New York: Columbia University Press.

Myerson, R. (1981): "Optimal Auction Design," *Mathematics of Operations Research*, 6, 58–78.

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

Pai, M. M. and R. Vohra (2013): "Optimal Dynamic Auctions and Simple Index Rules," *Mathematics of Operations Research*, 38, 682–697.

Parkes, D. C. and S. Singh (2003): "An MDP-Based Approach to Online Mechanism Design," in *Proceedings of the 17th Annual Conference on Neural Information Processing Systems (NIPS 03)*.

Pavan, A., I. Segal, and J. Toikka (2014): "Dynamic Mechanism Design: A Myersonian Approach," *Econometrica*, 82, 601–653.

Puterman, M. L. (1994): *Markov Decision Processes: Discrete Stochastic Dynamic Programming*, New York: Wiley.

Rochet, J.-C. and J. Tirole (2006): "Two-Sided Markets: A Progress Report," *RAND Journal of Economics*, 35, 645–667.

Rostek, M. and M. Weretka (2015): "Dynamic Thin Markets," *Review of Financial Studies*, 28, 2946–2992.

Satterthwaite, M. and A. Shneyerov (2007): "Dynamic Matching,Two-sided Incomplete Information, and Participation Costs: Existence and Convergence to Perfect Competition," *Econometrica*, 75, 155–200.

Satterthwaite, M. and S. R. Williams (2002): "The Optimality of a Simple Market Mechanism," *Econometrica*, 70, 1841–1863.

Skrzypacz, A. and J. Toikka (2015): "Mechanisms for Repeated Trade," *The American Economic Journal: Microeconomics*, 7, 252–293.

Stigler, G. (1966): *The Theory of Price*, New York: Macmillan, 3rd ed.

Ünver, M. U. (2010): "Dynamic Kidney Exchange," *Review of Economic Studies*, 77, 372–414.

Vayanos, D. (1999): "Strategic Trading and Welfare in a Dynamic Market," *Review of Economic Studies*, 66, 219–254.

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

Yavaş, A. (1996): "Matching of Buyers and Sellers by Brokers: A Comparison of Alternative

Commission Structures," *Real Estate Economics*, 24, 97–112.

Appendix

A Proofs

A.1 Proof of Proposition 1

Proof. We start by proving the result with Bayesian incentive compatibility and individual rationality constraints. Let V^{B_t} and C^{S_t} denote the respective types of the buyer and seller that arrive in period t. When B_t reports $\hat{v}^{B_t}$ and S_t reports $\hat{c}^{S_t}$ the respective interim discounted allocation probabilities, assuming all other agents arriving after period $t-1$ report truthfully, are given by

$$\begin{aligned} q^{B_t}(\hat{v}^{B_t}) &= \sum_{i=t}^{\infty} \sum_{\boldsymbol{h}_i \in \mathcal{H}_i} \delta^{i-1} Q_i^{B_t}(\boldsymbol{h}_i) \mathbb{P}(\boldsymbol{H}_i = \boldsymbol{h}_i | V^{B_t} = \hat{v}^{B_t}), \\ q^{S_t}(\hat{c}^{S_t}) &= \sum_{i=t}^{\infty} \sum_{\boldsymbol{h}_i \in \mathcal{H}_i} \delta^{i-1} Q_i^{S_t}(\boldsymbol{h}_i) \mathbb{P}(\boldsymbol{H}_i = \boldsymbol{h}_i | C^{S_t} = \hat{c}^{S_t}), \end{aligned} \tag{18}$$

where $\mathbb{P}(\boldsymbol{H}_i = \boldsymbol{h}_i | V^t = \hat{v}^{B_t})$ denotes the conditional probability that the period $i \geq t$ report history is $\boldsymbol{h}_i$, given that B_t reports $\hat{v}^{B_t}$ in period t, and analogously for $\mathbb{P}(\boldsymbol{H}_i = \boldsymbol{h}_i | C^{S_t} = \hat{c}^{S_t})$. Similarly, when B_t reports $\hat{v}^{B_t}$ and S_t reports $\hat{c}^{S_t}$ the respective expected interim discounted payments, assuming all other agents arriving after period $t-1$ report truthfully, are given by

$$\begin{aligned} m^{B_t}(\hat{v}^{B_t}) &= \sum_{i=t}^{\infty} \sum_{\boldsymbol{h}_i \in \mathcal{H}_i} \delta^{i-1} M_i^{B_t}(\boldsymbol{h}_i) \mathbb{P}(\boldsymbol{H}_i = \boldsymbol{h}_i | V^{B_t} = \hat{v}^{B_t}), \\ m^{S_t}(\hat{c}^{S_t}) &= \sum_{i=t}^{\infty} \sum_{\boldsymbol{h}_i \in \mathcal{H}_i} \delta^{i-1} M_i^{S_t}(\boldsymbol{h}_i) \mathbb{P}(\boldsymbol{H}_i = \boldsymbol{h}_i | C^{S_t} = \hat{c}^{S_t}). \end{aligned}$$

Next, we need to use the binding incentive compatibility and individual rationality constraints to write the payments in terms of the allocation rule and the agent types. First, for the worst-off types the individual rationality constraints bind and we have

$$m^{B_t}(v_1) = v_1 q^{B_t}(v_1) \quad \text{and} \quad m^{S_t}(c_m) = c_m q^{S_t}(c_m).$$

For all other buyer types $i \in \{2, \dots, n\}$, the incentive compatibility constraint binds locally downwards and we have

$$m^{B_t}(v_i) = v_i(q^{B_t}(v_i) - q^{B_t}(v_{i-1})) + m^{B_t}(v_{i-1}).$$

Setting $q^{B_t}(v_0) = 0$ for notational convenience and using recursive substitution we have, for all types $i \in \{1, \ldots, n\}$,

$$m^{B_t}(v_i) = \sum_{k=1}^{i} v_k(q^{B_t}(v_k) - q^{B_t}(v_{k-1})). \tag{19}$$

Similarly, for all other seller types $j \in \{1, \ldots, m-1\}$, the incentive compatibility constraint binds locally upwards and we have

$$m^{S_t}(c_j) = c_j(q^{S_t}(c_j) - q^{S_t}(c_{j+1})) + m^{S_t}(c_{j+1}).$$

Setting $q^{S_t}(c_{m+1}) = 0$ for notational convenience and using recursive substitution we have, for all types $j \in \{1, \ldots, m\}$,

$$m^{S_t}(c_j) = \sum_{k=j}^{m} c_k(q^{S_t}(c_k) - q^{S_t}(c_{k+1})). \tag{20}$$

Finally, starting from (4) can rewrite the virtual type functions as

$$\begin{aligned}
\Phi(v_i) &= v_i \left(1 + \frac{1 - F(v_i)}{f(v_i)}\right) - v_{i+1}\frac{1 - F(v_i)}{f(v_i)} \\
&= v_i \frac{1 - F(v_i) + f(v_i)}{f(v_i)} - v_{i+1}\frac{1 - F(v_i)}{f(v_i)} \\
&= v_i \frac{1 - F(v_{i-1})}{f(v_i)} - v_{i+1}\frac{1 - F(v_i)}{f(v_i)}
\end{aligned} \tag{21}$$

and

$$\begin{aligned}
\Gamma(c_j) &= c_j \left(1 + \frac{G(c_{j-1})}{g(c_j)}\right) - c_{j-1}\frac{G(c_{j-1})}{g(c_j)} \\
&= c_j \frac{G(c_{j-1}) + g(c_j)}{g(c_j)} - c_{j-1}\frac{G(c_{j-1})}{g(c_j)} \\
&= c_j \frac{G(c_j)}{g(c_j)} - c_{j-1}\frac{G(c_{j-1})}{g(c_j)}.
\end{aligned} \tag{22}$$

We can now write the profit of the designer in terms of agents' virtual types as follows

$$\begin{aligned}
R &= \sum_{i=1}^{\infty} \boldsymbol{E}\left[m^{B_i}(V^{B_i}) - m^{S_i}(C^{S_i})\right] \\
&= \sum_{i=1}^{\infty}\sum_{k=1}^{n} m^{B_i}(v_k)f(v_k) - \sum_{i=1}^{\infty}\sum_{k=1}^{m} m^{S_i}(c_k)g(c_k) \\
&= \sum_{i=1}^{\infty}\sum_{k=1}^{n}\sum_{\ell=1}^{k} v_\ell(q^{B_i}(v_\ell) - q^{B_i}(v_{\ell-1}))f(v_k) - \sum_{i=1}^{\infty}\sum_{k=1}^{m}\sum_{\ell=k}^{m} c_\ell(q^{S_i}(c_\ell) - q^{S_i}(c_{\ell+1}))g(c_k) \\
&= \sum_{i=1}^{\infty}\sum_{\ell=1}^{n}\sum_{k=\ell}^{n} v_\ell(q^{B_i}(v_\ell) - q^{B_i}(v_{\ell-1}))f(v_k) - \sum_{i=1}^{\infty}\sum_{\ell=1}^{m}\sum_{k=1}^{\ell} c_\ell(q^{S_i}(c_\ell) - q^{S_i}(c_{\ell+1}))g(c_k) \\
&= \sum_{i=1}^{\infty}\sum_{\ell=1}^{n} v_\ell(q^{B_i}(v_\ell) - q^{B_i}(v_{\ell-1}))(1 - F(v_{\ell-1})) - \sum_{i=1}^{\infty}\sum_{\ell=1}^{m} c_\ell(q^{S_i}(c_\ell) - q^{S_i}(c_{\ell+1}))G(c_\ell) \\
&= \sum_{i=1}^{\infty}\sum_{\ell=1}^{n}[v_\ell(1 - F(v_{\ell-1})) - v_{\ell+1}(1 - F(v_\ell))]q^{B_i}(v_\ell) \\
&- \sum_{i=1}^{\infty}\sum_{\ell=1}^{m}[c_\ell G(c_\ell) - c_{\ell-1}G(c_{\ell-1})]q^{S_i}(c_\ell) \\
&= \sum_{i=1}^{\infty}\sum_{\ell=1}^{n} \Phi(v_\ell)q^{B_i}(v_\ell)f(v_\ell) - \sum_{i=1}^{\infty}\sum_{\ell=1}^{m} \Gamma(c_\ell)q^{S_i}(c_\ell)g(c_\ell).
\end{aligned}$$

Here, the first line used the definition of designer profit, the second line computed the expectation over the types of the period i buyer and seller, line three used (19) and (20), line four interchanged the order of the inner summations, line five computed the innermost summation, line six collected terms by allocation rather than type and line seven used (21) and (22). Next, using (18) gives

$$\begin{aligned}
R &= \sum_{i=1}^{\infty}\sum_{\ell=1}^{n} \Phi(v_\ell) \sum_{t=i}^{\infty}\sum_{\boldsymbol{h}_t \in \mathcal{H}_t} \delta^{t-1} Q_t^{S_i} \mathbb{P}(\boldsymbol{H}_t = \boldsymbol{h}_t \mid V^{B_i} = v_\ell) f(v_\ell) \\
&- \sum_{i=1}^{\infty}\sum_{\ell=1}^{m} \Gamma(c_\ell) \sum_{t=i}^{\infty}\sum_{\boldsymbol{h}_t \in \mathcal{H}_t} \delta^{t-1} Q_t^{B_i} \mathbb{P}(\boldsymbol{H}_t = \boldsymbol{h}_t \mid C^{S_i} = c_\ell) g(c_\ell) \\
&= \sum_{i=1}^{\infty}\sum_{t=i}^{\infty}\sum_{\boldsymbol{h}_t \in \mathcal{H}_t}\sum_{\ell=1}^{n} \delta^{t-1}\Phi(v_\ell) Q_t^{B_i}(\boldsymbol{h}_t) \mathbb{P}(\boldsymbol{H}_t = \boldsymbol{h}_t | V^{B_i} = v_\ell) f(v_\ell) \\
&- \sum_{i=1}^{\infty}\sum_{t=i}^{\infty}\sum_{\boldsymbol{h}_t \in \mathcal{H}_t}\sum_{\ell=1}^{m} \delta^{t-1}\Gamma(c_\ell) Q_t^{S_i}(\boldsymbol{h}_t) \mathbb{P}(\boldsymbol{H}_t = \boldsymbol{h}_t | C^{S_i} = c_\ell) g(c_\ell)
\end{aligned}$$

and using the law of total probability we have

$$\begin{aligned}R &= \sum_{i=1}^{\infty}\sum_{t=i}^{\infty}\sum_{\boldsymbol{h}_t\in\mathcal{H}_t}\delta^{t-1}\Phi(v^{B_i}(\boldsymbol{h}_t))Q_t^{B_i}(\boldsymbol{h}_t)\mathbb{P}(\boldsymbol{H}_t=\boldsymbol{h}_t)\\ &-\sum_{i=1}^{\infty}\sum_{t=i}^{\infty}\sum_{\boldsymbol{h}_t\in\mathcal{H}_t}\delta^{t-1}\Gamma(c^{S_i}(\boldsymbol{h}_t))Q_t^{S_i}(\boldsymbol{h}_t)\mathbb{P}(\boldsymbol{H}_t=\boldsymbol{h}_t)\\ &=\sum_{t=1}^{\infty}\sum_{i=1}^{t}\sum_{\boldsymbol{h}_t\in\mathcal{H}_t}\delta^{t-1}\left(\Phi(v^{B_i}(\boldsymbol{h}_t))Q_t^{B_i}(\boldsymbol{h}_t)-\Gamma(c^{S_i}(\boldsymbol{h}_t))Q_t^{S_i}(\boldsymbol{h}_t)\right)\mathbb{P}(\boldsymbol{H}_t=\boldsymbol{h}_t)\end{aligned}$$

as required.

Repeating this procedure, we can show that the result also holds under interim and periodic ex post incentive constraints. Under interim incentive constraints, we let $m(\hat{\theta},\boldsymbol{h}_{t-1})$ denote the expected discounted payment for an agent that reports $\hat{\theta}$ at history $\boldsymbol{h}_{t-1}$ and compute

$$\begin{aligned}R &= \sum_{i=1}^{\infty}\boldsymbol{E}_{V_i,C_i,\boldsymbol{H}_{t-1}}\left[m^{B_i}(V_i,\boldsymbol{H}_{t-1})-m^{S_i}(C_i,\boldsymbol{H}_{t-1})\right]\\ &=\sum_{i=1}^{\infty}\sum_{\boldsymbol{h}_{t-1}\in\mathcal{H}_{t-1}}\left[\sum_{\ell=1}^{n}\Phi(v_\ell)q^{B_i}(v_\ell,\boldsymbol{h}_{t-1})f(v_\ell)-\sum_{\ell=1}^{m}\Gamma(c_\ell)q^{S_i}(c_\ell,\boldsymbol{h}_{t-1})g(c_\ell)\right]\mathbb{P}(\boldsymbol{H}_{t-1}=\boldsymbol{h}_{t-1})\\ &=\sum_{i=1}^{\infty}\sum_{t=i}^{\infty}\sum_{\boldsymbol{h}_t\in\mathcal{H}_t}\sum_{\boldsymbol{h}_{t-1}\in\mathcal{H}_{t-1}}\delta^{t-1}\left[\sum_{\ell=1}^{n}\Phi(v_\ell)Q_t^{B_i}(\boldsymbol{h}_t)\mathbb{P}(\boldsymbol{H}_t=\boldsymbol{h}_t|V^{B_i}=v_\ell,\boldsymbol{H}_{t-1}=\boldsymbol{h}_{t-1})f(v_\ell)\right.\\ &\left.-\sum_{\ell=1}^{m}\Gamma(c_\ell)Q_t^{S_i}(\boldsymbol{h}_t)\mathbb{P}(\boldsymbol{H}_t=\boldsymbol{h}_t|C^{S_i}=c_\ell,\boldsymbol{H}_{t-1}=\boldsymbol{h}_{t-1})g(c_\ell)\right]\mathbb{P}(\boldsymbol{H}_{t-1}=\boldsymbol{h}_{t-1}).\end{aligned}$$

Applying the law of total probability we then have

$$R=\sum_{t=1}^{\infty}\sum_{i=1}^{t}\sum_{\boldsymbol{h}_t\in\mathcal{H}_t}\delta^{t-1}\left(\Phi(v^{B_i}(\boldsymbol{h}_t))Q_t^{B_i}(\boldsymbol{h}_t)-\Gamma(c^{S_i}(\boldsymbol{h}_t))Q_t^{S_i}(\boldsymbol{h}_t)\right)\mathbb{P}(\boldsymbol{H}_t=\boldsymbol{h}_t)$$

as required. Intuitively, conditioning period t allocations and payments for each agent on the period $t-1$ history makes no difference to the profit of the designer when we sum over all histories.

Similarly, under periodic ex post incentive constraints, we let $m(\hat{\theta},\vartheta,\boldsymbol{h}_{t-1})$ denote the expected discounted payment for an agent that reports $\hat{\theta}$ at history $\boldsymbol{h}_{t-1}$ when the other period t agent reports ϑ and compute

$$
\begin{aligned}
R &= \sum_{i=1}^{\infty} \boldsymbol{E}_{V_i, C_i, \boldsymbol{H}_{t-1}} \left[m^{B_i}(V_i, C_i, \boldsymbol{H}_{t-1}) - m^{S_i}(C_i, V_i, \boldsymbol{H}_{t-1}) \right] \\
&= \sum_{i=1}^{\infty} \sum_{\boldsymbol{h}_{t-1} \in \mathcal{H}_{t-1}} \sum_{\ell=1}^{n} \sum_{k=1}^{m} \left[\Phi(v_\ell) q^{B_i}(v_\ell, c_k, \boldsymbol{h}_{t-1}) - \Gamma(c_k) q^{S_i}(c_k, v_\ell, \boldsymbol{h}_{t-1}) \right] f(v_\ell) g(c_\ell) \\
&\times \mathbb{P}(\boldsymbol{H}_{t-1} = \boldsymbol{h}_{t-1}) \\
&= \sum_{i=1}^{\infty} \sum_{t=i}^{\infty} \sum_{\boldsymbol{h}_t \in \mathcal{H}_t} \sum_{\boldsymbol{h}_{t-1} \in \mathcal{H}_{t-1}} \sum_{\ell=1}^{n} \sum_{k=1}^{m} \delta^{t-1} \left[\Phi(v_\ell) Q_t^{B_i}(\boldsymbol{h}_t) - \Gamma(c_k) Q_t^{S_i}(\boldsymbol{h}_t) \right] \\
&\times \mathbb{P}(\boldsymbol{H}_t = \boldsymbol{h}_t | V^{B_i} = v_\ell, C^{S_i} = c_k, \boldsymbol{H}_{t-1} = \boldsymbol{h}_{t-1}) f(v_\ell) g(c_k) \mathbb{P}(\boldsymbol{H}_{t-1} = \boldsymbol{h}_{t-1}).
\end{aligned}
$$

Applying the law of total probability we then have

$$
R = \sum_{t=1}^{\infty} \sum_{i=1}^{t} \sum_{\boldsymbol{h}_t \in \mathcal{H}_t} \delta^{t-1} \left(\Phi(v^{B_i}(\boldsymbol{h}_t)) Q_t^{B_i}(\boldsymbol{h}_t) - \Gamma(c^{S_i}(\boldsymbol{h}_t)) Q_t^{S_i}(\boldsymbol{h}_t) \right) \mathbb{P}(\boldsymbol{H}_t = \boldsymbol{h}_t)
$$

as required. □

A.2 Proof of Theorem 1

Proof. The proof proceeds by deriving the optimal policy of the Markov decision process and mapping this to the optimal allocation rule, before checking the relevant incentive constraints.

First, note that by Theorem 6.2.10 of Puterman (1994) there exists a deterministic stationary optimal policy of the Markov decision process. Second, note that the optimal policy must immediately clear all efficient pairs. Third, note that sample paths of the Markov decision process are such that if x_S suboptimal trades are stored in a given period, $x_S - 1$ trades must have been stored in some previous period. Thus, if x_S suboptimal trades are stored under the stationary optimal policy, it must be optimal to retain $x_S - 1$ trades.

An unbounded number of suboptimal pairs cannot be stored under the optimal policy since as the number of stored suboptimal pairs diverges to infinity, the expected number of periods until an additional stored suboptimal pair is rematched diverges to infinity. Thus, the expected discounted benefit of storing an additional suboptimal pair converges to zero, while the benefit of immediately clearing a suboptimal pair is always $\Delta_\alpha > 0$. Putting everything together, there exists a maximum number τ^* of suboptimal trades which can be optimally stored and the optimal policy π^* is a threshold policy.

For the incentive constraints, it suffices to show that $q^{B_t}(\overline{v}) \geq q^{B_t}(\underline{v})$ and $q^{S_t}(\underline{c}) \geq q^{S_t}(\overline{c})$. The allocation rule induced by the optimal threshold policy π^* is unique up to the queueing protocol for storing suboptimal pairs. We let $q(\theta, \vartheta)$ denote the expected discounted probability of trade under the optimal policy for an agent of type θ, who arrives with an agent of type ϑ.[37] Start by considering the arrival of a buyer of type $\overline{v}$ in period t. If this buyer is paired with a seller of type $\underline{c}$ (which occurs with probability p), that trade is immediately executed. Otherwise, the buyer will be stored as part of a suboptimal pair. We have

$$q^{B_t}(\overline{v}) = \delta^{t-1}\left[p + (1-p)q(\overline{v}, \overline{c})\right]. \tag{23}$$

Next, consider the arrival of a buyer of type $\underline{v}$ in period t. This agent trades with non-zero probability only if it arrives as part of a suboptimal pair and we have

$$q^{B_t}(\underline{v}) = \delta^{t-1} p q(\underline{v}, \underline{c}). \tag{24}$$

Comparing (23) and (24) we see that $q^{B_t}(\overline{v}) \geq q^{B_t}(\underline{v})$ since $q(\underline{v}, \underline{c}) \leq 1$. An analogous argument shows that $q^{S_t}(\underline{c}) \geq q^{S_t}(\overline{c})$. □

A.3 Proof of Corollary 1

Proof. The first part of Corollary 1 follows immediately from the fact that inefficient $(\underline{v}, \overline{c})$ pairs neither contribute to social welfare nor the profit of the designer. We now prove the second part of the corollary. We start by noting that by definition the designer's objective function is maximized under any optimal market clearing policy π^* and so $(1-\alpha)W + \alpha R$ is invariant to the queueing protocol. Furthermore, the queueing protocol cannot affect social welfare W since it is only used to break ties among traders of the same type. Thus, the profit of the designer R is also invariant to the queueing protocol. □

[37] Here, we take the expectation over the types of past and future agents. We can be agnostic about the queueing protocol for suboptimal pairs.

A.4 Proof of Proposition 2

Proof. The transition matrix $\boldsymbol{P}$ of the order book Markov chain $\{Y_t\}_{t\in\mathbb{N}}$ is given by

$$\boldsymbol{P} = \begin{pmatrix} 1-2\lambda & 2\lambda & 0 & \cdots & 0 & 0 & 0 \\ \lambda & 1-2\lambda & \lambda & & 0 & 0 & 0 \\ 0 & \lambda & 1-2\lambda & & 0 & 0 & 0 \\ \vdots & \ddots & \ddots & \ddots & \ddots & \ddots & \vdots \\ 0 & 0 & 0 & & 1-2\lambda & \lambda & 0 \\ 0 & 0 & 0 & & \lambda & 1-2\lambda & \lambda \\ 0 & 0 & 0 & \cdots & 0 & \lambda & 1-\lambda \end{pmatrix},$$

where $\lambda = p(1-p)$. so we are dealing with a simple birth-and-death process (see, for example, pages 184–189 in Borovkov (2014)). For such processes, it is well-known that the stationary distribution $\boldsymbol{\kappa}$ of the Markov chain $\{Y_t\}_{t\in\mathbb{N}}$ is given by

$$\kappa_0 = \frac{1}{2\tau+1}, \quad \forall i \in \{1,\dots,\tau\}, \quad \kappa_i = \frac{2}{2\tau+1}.$$

We now compute the market maker's expected period t payoff when the market is stationary under the threshold policy with threshold τ. With probability p^2 a $(\overline{v},\underline{c})$ pair arrives and with probability $(1-p)^2$ a $(\overline{v},\underline{c})$ pair arrives, creating respective payoffs of 1 and 0. A suboptimal pair arrives with probability $2p(1-p)$. With probability $\tau/(2\tau+1)$ this pair arrives to a market in which non-identical suboptimal pairs are stored. In this case, it is rematched to create an efficient trade which is immediately cleared. With probability $\tau/(2\tau+1)$ the number of identical suboptimal pairs stored is less than τ and the arriving pair is stored. Finally, with probability $1/(2\tau+1)$ the maximum number of identical suboptimal pairs are stored and the one suboptimal pair is immediately cleared. Thus, assuming the market is stationary, the market maker's expected period t payoff is

$$p^2 + \frac{2p(1-p)(\Delta_\alpha+\tau)}{2\tau+1}.$$

$\square$

A.5 Proof of Proposition 3

Proof. Since the optimal policy is stationary, the first part of the proposition can be proven using the stationary payoff. If the designer changes the market clearing threshold from τ to

$\tau+1$, the expected change in welfare under the stationary distribution is

$$W^D(\tau+1)-W^D(\tau)=\frac{1}{1-\delta}\frac{2p(1-p)(1-2\Delta_\alpha)}{(2\tau+3)(2\tau+1)}.$$

Differentiating with respect to the problem parameters, we obtain

$$\begin{aligned}
\frac{\partial(W^D(\tau+1)-W^D(\tau))}{\partial\delta} &= \frac{1}{(1-\delta)^2}\frac{2p(1-p)(1-2\Delta_\alpha)}{(2\tau+3)(2\tau+1)} > 0, \\
\frac{\partial(W^D(\tau+1)-W^D(\tau))}{\partial p(1-p)} &= \frac{1}{1-\delta}\frac{2(1-2\Delta_\alpha)}{(2\tau+3)(2\tau+1)} > 0 \\
\frac{\partial(W^D(\tau+1)-W^D(\tau))}{\partial\Delta_\alpha} &= -\frac{1}{1-\delta}\frac{4p(1-p)}{(2\tau+3)(2\tau+1)} < 0.
\end{aligned}$$

Since the payoff associated with increasing τ is increasing in $p(1-p)$ and decreasing in Δ_α, so too is τ^*.

Next, examining (13) reveals that for $x_S \in \{0,1,\dots,\tau^*\}$ an increase in Δ_α leads to an increase in $V^D_{\tau^*}(x_S)$. Since the total expected discounted payoff is increasing for each state, total expected discounted welfare gains are increasing in Δ_α. For $x_S \in \{1,\dots,\tau^*-1\}$, ranking the outcomes on the right-hand side of (13) by payoff gives $1+V^D_{\tau^*}(x_S) > 1+V^D_{\tau^*}(x_S-1) > V^D_{\tau^*}(x_S+1) > V^D_{\tau^*}(x_S)$. The outcomes $1+V^D_{\tau^*}(x_S-1)$ and $V^D_{\tau^*}(x_S+1)$ occur with equal probability and an increase in p leads to an increase in the probability of the best outcome and a decrease in the probability of the worst outcome. Since similar reasoning applies to the boundary equations (that is, those corresponding to $x_S=0$ and $x_S=\tau^*$), an increase in p increases the total expected discounted payoff for each state. Thus, total expected welfare is increasing in p.

Finally, noting that Δ_α is increasing Δ_0, we see that all statements regarding Δ_α also hold for Δ_0. □

A.6 Proof of Corollary 2

Proof. Differentiating Δ_α yields

$$\frac{\partial\Delta_\alpha}{\partial\alpha} = -\frac{p}{1-p}(1-\Delta_0) < 0.$$

Recall from the proof of Proposition 3 that τ^* is decreasing in Δ_α. Combining this with the previous result shows that τ^* is increasing in α. Thus, market thickness as measured by τ^* is increasing in α. □

A.7 Proof of Proposition 4

Proof. Suppose that $\tau^* > 0$ for $\alpha = 0$. We have already argued that the balanced budget posted price mechanism with τ^* implements the efficient allocation rule, assuming truthful reporting. Furthermore, the balanced budget posted price mechanism does not run a deficit by construction. However, we need to check the P-IC and P-IR incentive constraints. First, note that agents of type $\underline{v}$ and $\overline{c}$ receive a payoff of zero whenever they report truthfully, since under truthful reporting these agents will only accept prices of $\underline{v}$ and $\overline{c}$ respectively. This holds regardless of the history and the types of contemporary agents. If these agents do not report truthfully and accept a posted price of $1/2$, they will receive a negative expected discounted payoff. Again, this holds regardless of the history and the types of contemporary agents. Thus, the P-IC and P-IR constraints are satisfied. If agents of type $\overline{v}$ are offered a posted price of Δ_0, they will clearly accept regardless of the history and of the types of contemporary agents. If a price of $1/2$ or $1 - \Delta_0$ is observed, misreporting guarantees that the buyer will eventually leave the market without trading (either immediately or later due to the last-come-first-served queueing protocol), regardless of the history and the types of contemporary agents. However, reporting truthfully ensures that the buyer will eventually trade either at a price of $1/2$ or $1-\Delta_0$. Thus, the B-IC constraint holds for buyers of type $\overline{v}$. Similarly, it is clear that truthfully reporting guarantees buyers of type $\overline{v}$ a positive expected discounted payoff regardless of the history and regardless of the types of contemporary agents. Therefore, the P-IR constraint is also satisfied. Since a similar argument applies to sellers of type $\underline{c}$, we have a P-IC and P-IR mechanism as required.

Next, suppose that the efficient allocation rule can be implemented using a P-IC and P-IR budget balanced posted price mechanism. Then we must have $\tau^* > 0$ for $\alpha = 0$, since the balanced budget posted price mechanism cannot implement the efficient allocation rule if $\tau^* = 0$. This follows immediately from the fact that there is no price that can clear efficient trades and both types of suboptimal trades. □

A.8 Proof of Corollary 3

Proof. For $\alpha = 0$, both the budget balanced posted price mechanism and the optimal mechanism implement the efficient allocation. However, the optimal mechanism maximizes the profit of the market maker within the class of efficient, P-IC and P-IR mechanisms and is not subject to the posted price constraint. Therefore, since the balanced budget mechanism does not run a deficit, neither can the optimal mechanism. □

A.9 Proof of Proposition 5

Note that as this proof pertains to the general setup from Section 2.1, we use the Markov decision process notation from Online Appendix B.1. We also let $V_{\pi}^{D}(\boldsymbol{x})$ denote the expected discounted present value of being in state $\boldsymbol{x}$ under the policy π. Before proving the proposition, we prove the following lemma.

Lemma A1. *Let the optimal policy π^* be given and consider the associated order book Markov chain of positive recurrent states $\mathcal{Y}^*$. Then for any policies π and π' and associated Markov chains $\mathcal{Y}$ and $\mathcal{Y}'$ such that $\mathcal{Y} \subset \mathcal{Y}' \subset \mathcal{Y}^*$ we have $V_{\pi^*}^{D}(\boldsymbol{x}) \geq V_{\pi'}^{D}(\boldsymbol{x}) \geq V_{\pi}^{D}(\boldsymbol{x})$ for all states $\boldsymbol{x} \in \mathcal{X}$.*

Proof. That $V_{\pi^*}^{D}(\boldsymbol{x}) \geq V_{\pi'}^{D}(\boldsymbol{x})$ and $V_{\pi^*}^{D}(\boldsymbol{x}) \geq V_{\pi}^{D}(\boldsymbol{x})$ for all states $\boldsymbol{x} \in \mathcal{X}$ follows from the principle of optimality of dynamic programming. That $V_{\pi'}^{D}(\boldsymbol{x}) \geq V_{\pi}^{D}(\boldsymbol{x})$ follows from the fact that $\mathcal{Y}' \setminus \mathcal{Y} \subset \mathcal{Y}^*$, so storing in the states $\mathcal{Y}' \setminus \mathcal{Y}$ is optimal under π^*. □

Proof. First, we consider the continuum limit we obtain as $\delta \to 1$. With discrete type spaces there exists an interval $[\underline{p}^W, \overline{p}^W]$ of prices that implement the Walrasian allocation (with $\underline{p}^W \neq \overline{p}^W$ provided we exclude trades that generate zero social surplus). Thus, in the $\delta \to 1$ limit the optimal mechanism with $\alpha = 0$ can run a surplus in every period simply by posting a price of $p^B = \overline{p}^W$ for buyers and a price of $p^S = \underline{p}^W$ for sellers and the profit of the designer diverges. Next, let $\mathcal{Y}^{D,0}(\delta)$ and $\mathcal{Y}^{D,1}(\delta)$ denote the set of positive recurrent states of the order book Markov chain under the optimal policies for welfare-maximizing and profit-maximizing discriminatory market clearing, respectively. Recall that these functions are both increasing in δ. Under the assumption of dynamic regularity we have $\mathcal{Y}^{D,0}(\delta) \subset \mathcal{Y}^{D,1}(\delta)$ by Proposition B2. It follows that profit $R^{D,0}(\delta)$ under welfare maximization increases discontinuously at points at which the set $\mathcal{Y}^{D,0}(\delta)$ increases by Lemma A1 but may decrease in δ at points at which $\mathcal{Y}^{D,0}(\delta)$ is constant. However, since $\liminf_{\delta \to 1} R^{D,0}(\delta) = \infty$ it follows that there exists a

$\overline{\delta} < 1$ such that $R^{D,0}(\delta) > 0$ for all $\delta > \overline{\delta}$. That the optimal policies are implementable using a P-IC and P-IR mechanism follows from Theorem B1 (see Online Appendix B.1). □

A.10 Proof of Proposition 6

Proof. Clearly, the efficient allocation cannot be implemented by a posted price mechanism with a larger bid-ask spread, so we only need to check that the relative incentive constraints hold. First, agents of type $\underline{v}$ and $\overline{c}$ can guarantee themselves a payoff of 0 by always accepting respective prices of $p_B = \Delta_0$ and $p_S = 1 - \Delta_0$ and rejecting otherwise. Thus, the P-IR constraints hold for these types. Furthermore, these agents do strictly worse if they accept prices of $p_B = 1$ or $p_S = 0$ because they then receive a negative expected payoff. Thus, the P-IC constraints for worst-off types hold. The P-IR constraints hold for the efficient types because these types guarantee themselves a non-negative expected discounted payoff by always accepting every posted price. Thus, to complete the proof we only need to verify that buyers of type $\overline{v}$ will accept a posted price of $p_B = 1$ whenever they arrive to an order book below capacity (the argument for sellers of type $\underline{c}$ is analogous). Clearly, if $(\overline{v}, \overline{c})$ pairs are stored, buyers of type $\overline{v}$ cannot benefit from rejecting a price of $p_B = 1$ because this will result in their removal from the market. Therefore, we suppose that $\iota \in \{0, 1, \ldots, \tau^* - 1\}$ $(\underline{v}, \underline{c})$ suboptimal pairs are stored, so that rejecting $p_B = 1$ either results in the buyer being stored as part of a suboptimal pair or removed from the market. A buyer of type $\overline{v}$ can only benefit from rejecting a price of $p_B = 1$ if they are stored as part of a suboptimal trade and they then trade with non-zero probability at a price of $p_B = \Delta_0$. However, this can never happen with a last-come-first-serve protocol because once τ^* pairs are stored and the designer posts a price of $p_B = \Delta_0$, that price will always be accepted by the arriving buyer who has priority over the stored buyers. Thus, all of the appropriate incentive constraints have been verified. □

A.11 Proof of Theorem 2

Proof. Under discriminatory, uniform and fixed frequency market clearing per period welfare converges to p as $\delta \to 1$ and under instantaneous market clearing is always given by $p^2 + 2p(1-p)\Delta_0$. Hence, $\lim_{\delta \to 1} W^{U,\alpha}(\delta)/W^{D,\alpha}(\delta) = \lim_{\delta \to 1} W^{F,\alpha}(\delta)/W^{U,\alpha}(\delta) = 1$ and

$\lim_{\delta\to 1} W^{0,\alpha}(\delta)/W^{F,\alpha}(\delta) = p + 2(1-p)\Delta_0$ and the desired result immediately follows. □

A.12 Proof of Theorem 3

For ease of exposition, we formally prove the theorem for the $k = D$ case. The other cases are similar. However, before proceeding to the proof of the main result, we start by stating and proving the following lemma (which applies specifically to discriminatory market clearing, a more general version can be found in Lemma A1).

Lemma A2. *Let the optimal threshold τ^* under discriminatory market clearing be given. Fix a threshold τ and take $i \in \{0, 1, \ldots, \tau\}$. Then $V^D_\tau(i)$ increases in τ for $\tau \in \{0, 1, \ldots, \tau^*-1\}$, decreases in τ for $\tau \geq \tau^* + 1$ and attains a global maximum at $\tau = \tau^*$.*

Proof. First, take $\tau = \tau^* + 1$. Then $V^D_{\tau^*}(i) \geq V^D_{\tau^*+1}(i)$ by the principle of optimality of dynamic programming. For $\tau = \tau^* + 2$ we have $V^D_\tau(\tau) \leq r(0,1) + V^D_{\tau-1}(\tau) \leq r(0,1) + V^D_{\tau-1}(\tau-1) \leq r(0,2) + V^D_{\tau^*}(\tau^*)$, where $r(a_E, a_S)$ is the reward from clearing a_E efficient and a_S suboptimal pairs. This follows from the fact that clearing two suboptimal pairs in state τ is optimal under τ^*, so clearing one suboptimal pair in state τ must yield a higher payoff than clearing none. Iterating, we have that $V^D_\tau(i)$ decreases in τ for $\tau \geq \tau^*$. Next, set $\tau = \tau^* - 1$. Then $V^D_{\tau^*}(i) \geq V^D_{\tau^*-1}(i)$ by the principle of optimality of dynamic programming. Next, setting $\tau = \tau^* - 2$ we have $r(0,1) + V^D_\tau(\tau) \leq r(0,1) + V^D_{\tau+1}(\tau) \leq V^D_{\tau+1}(\tau+1)$ since storing in state $\tau + 1$ is optimal under τ^*. Iterating, we have that $V^D_\tau(i)$ increases in τ for $\tau \in \{0, 1, \ldots, \tau^*\}$. Putting all of this together implies that we have a global maximum at $\tau = \tau^*$. □

We now prove Theorem 3 for $k = D$.

Proof. First, let $W^\infty(\delta)$ denote welfare under the market clearing policy with $\tau^* = \infty$ (that is, the mechanism under which all suboptimal trades are stored indefinitely and only efficient trades are executed). Further, let $\tau^{D,0}(\delta)$ and $\tau^{D,1}(\delta)$ denote the optimal thresholds under welfare-maximizing and profit-maximizing discriminatory market clearing, respectively. Recall that these functions are both increasing in δ. Further, $W^\infty(\delta)$ and $W^{0,0}(\delta)$ are continuous, increasing functions of δ and there exists $\tilde{\delta}$ such that $W^\infty(\tilde{\delta}) > W^{0,0}(\tilde{\delta})$ because

$\tau^{D,0}(\delta) \to \infty$ as $\delta \to 1$. Since $\tau^{D,0}(\delta) \leq \tau^{D,1}(\delta) < \infty$ for $\delta \in [0,1)$ by Corollary 2, Lemma A2 implies that $W^{D,1}(\delta) \geq W^{\infty}(\delta)$ for all $\delta \in [0,1)$. We also have that $W^{D,1}(\delta)$ decreases discontinuously at points at which $\tau^{D,1}(\delta)$ increases (again, by Lemma A2) and increases continuously in δ at all points at which $\tau^{D,1}(\delta)$ is constant (that is, the underlying market clearing policy does not vary). It immediately follows that there exists $\delta_D \leq \tilde{\delta}$ such that $W^{D,1}(\delta_D) = W^{0,0}(\delta_D)$ and $W^{D,1}(\delta) > W^{0,0}(\delta)$ for $\delta > \delta_D$. □

Online Appendices

B Markov decision process methodology

In Section 3 we used Markov decision process techniques to determine the allocation rules for the class of Bayesian optimal mechanisms with binary type spaces. In this appendix we illustrate the flexibility of this methodology by considering a variety of extensions. First, we characterize the class of Bayesian optimal mechanisms with the general discrete type spaces. We then consider various constraints imposed on the designer by deriving the optimal mechanisms under uniform market clearing and fixed frequency market clearing, which were introduced in Section 5.

B.1 General discrete types

We now show how the Markov decision process analysis from Section 3 generalizes to the type spaces $\mathcal{V} = \{v_1, \ldots, v_n\}$ and $\mathcal{C} = \{c_1, \ldots, c_m\}$ with $v_i < v_{i+1}$ and $c_j < c_{j+1}$ introduced in Section 2. The state space of the Markov decision process is now $\mathcal{X} = \mathbb{Z}_{\geq 0}^{n+m}$ with typical state $\boldsymbol{x} = (x_1^B, \ldots, x_n^B, x_1^S, \ldots, x_m^S)$, where x_i^B denotes the number of buyers that reported to be of type v_i and x_j^S denotes the number of sellers that reported to be of type c_j. We again let $\boldsymbol{X}_t \in \mathcal{X}$ denote the state of the market after the arrival of period t agents. Next, letting $\mathcal{A}_{\boldsymbol{x}}$ denote the set of actions available in state $\boldsymbol{x}$ we have $\mathcal{A}_{\boldsymbol{x}} = \{(a_1^B, \ldots, a_n^B, a_1^S, \ldots, a_m^S) \in \mathbb{Z}_{\geq 0} : a_i^B \leq x_i^B \ \forall i, \ a_j^S \leq x_j^S \ \forall j\}$. We again set $\mathcal{A} = \cup_{\boldsymbol{x} \in \mathcal{X}} \mathcal{A}_{\boldsymbol{x}}$ and let $\boldsymbol{A}_t$ denote the action taken by the designer in period $t \in \mathbb{N}$.

Let $\boldsymbol{e}_i^B \in \mathbb{Z}^{n+m}$ denote the unit vector with a one in the ith component and let $\boldsymbol{e}_j^S \in \mathbb{Z}^{n+m}$ denote the unit vector with a one in the $(n+j)$th component. Letting $P_{\boldsymbol{a}}(\boldsymbol{x}, \boldsymbol{x'})$ denote the probability that $\boldsymbol{X}_{t+1} = \boldsymbol{x'}$, given $\boldsymbol{X}_t = \boldsymbol{x}$ and feasible $\boldsymbol{A}_t = \boldsymbol{a}$, we have $P_{\boldsymbol{a}}(\boldsymbol{x}, \boldsymbol{x'}) = f(v_i)g(c_j)\mathbf{1}(\boldsymbol{x'} = \boldsymbol{x} - \boldsymbol{a} + \boldsymbol{e}_i^B + \boldsymbol{e}_j^S)$. It remains to define the reward function given some action $\boldsymbol{a}$. We first introduce the inverse demand function associated with action $\boldsymbol{a}$ which, for $i \in \{1, \ldots, n\}$ and $k \in \{\sum_{\ell=i+1}^n a_\ell^B + 1, \ldots, \sum_{\ell=i}^n a_\ell^B\}$, is given by

$$\mathbb{D}_{\boldsymbol{a}}(k) = \Phi_\alpha(v_i),$$

as well as the inverse supply function associated with action $\boldsymbol{a}$ which, for $i \in \{1, \ldots, m\}$ and

$k \in \{\sum_{\ell=1}^{i-1} a_\ell^S + 1, \ldots, \sum_{\ell=1}^{i} a_\ell^S\}$, is given by

$$\mathbb{S}_{\boldsymbol{a}}(k) = \Gamma_\alpha(c_i).$$

Letting $N = \min\left\{\sum_{i=1}^n a_i^B, \sum_{j=1}^m a_j^S\right\}$, the reward $r(\boldsymbol{a})$ associated with implementing action $\boldsymbol{a}$ is then given by

$$r(\boldsymbol{a}) = \sum_{k=1}^{N} (\mathbb{D}_{\boldsymbol{a}}(k) - \mathbb{S}_\alpha(k))\mathbf{1}(\mathbb{D}_\alpha(k) \geq \mathbb{S}_{\boldsymbol{a}}(k)).$$

Thus, we have a Markov decision process $\langle \mathcal{X}, \mathcal{A}, P, r, \delta \rangle$. Since the state space $\mathcal{X}$ is countable, the feasible action sets $\mathcal{A}_{\boldsymbol{x}}$ are finite for all states $\boldsymbol{x}$ and the reward function is deterministic, a stationary deterministic optimal policy exists and is characterized by the appropriate Bellman equation (see, for example, Theorems 6.2.6 and 6.2.10 of Puterman (1994)). In particular, we have

$$\begin{aligned} V^D(\boldsymbol{x}) &= \max_{\boldsymbol{a} \in \mathcal{A}_{\boldsymbol{x}}} \left\{ r(\boldsymbol{a}) + \delta \sum_{i=1}^{n} \sum_{j=1}^{n} P_{\boldsymbol{a}}(\boldsymbol{x'}, \boldsymbol{x} - \boldsymbol{a} + \boldsymbol{e}_i^B + \boldsymbol{e}_j^S) V^D(\boldsymbol{x} - \boldsymbol{a} + \boldsymbol{e}_i^B + \boldsymbol{e}_j^S) \right\} \\ &= r(\pi^*(\boldsymbol{x})) + \delta \sum_{i=1}^{n} \sum_{j=1}^{n} P_{\pi^*(\boldsymbol{x})}(\boldsymbol{x'}, \boldsymbol{x} - \pi^*(\boldsymbol{x}) + \boldsymbol{e}_i^B + \boldsymbol{e}_j^S) V^D(\boldsymbol{x} - \pi^*(\boldsymbol{x}) + \boldsymbol{e}_i^B + \boldsymbol{e}_j^S). \end{aligned}$$

The optimal policy π^* induces an order book Markov chain with positive recurrent states which we denote by $\mathcal{Y}$. There is an increase in market thickness under policy π relative to policy π' if $\mathcal{Y}(\pi') \subset \mathcal{Y}(\pi)$.

In general, it may be optimal to store infeasible trades. However, we assume that as soon as the state is such that a given agent will never trade under the optimal policy, that agent immediately leaves the market without trading. This assumption ensures that the order book Markov chain is stationary. For example, assuming that (v_1, c_m) trades are infeasible in the sense that $v_1 < c_m$ (which is needed so that the optimal policy does not merely stipulate that every pair is immediately cleared from the market) agents that form such pairs will be immediately cleared from the market. Provided there exist feasible trading partners for all other types, it may be optimal to store all other types of trading pairs in the order book.

Theorem B1. *The optimal market clearing policy can be implemented using a P-IC and P-IR mechanism.*

Proof. The allocation rule induced by the optimal market clearing policy π^* is unique up to the queueing protocol (which is a tie-breaking rule that does not affect the objective function $\alpha R + (1-\alpha)W$ of the designer). So fix any queueing protocol and the associated allocation rule. For the incentive constraints, it suffices to show that in any period t the allocation rules q^{B_t} and q^{S_t} are increasing and decreasing respectively. For any $i \in \{1, \ldots, n-1\}$ consider the arrival of a buyer of type v_i in period t and any sample path on which that buyer trades in period $T \geq t$. Now suppose we replace the buyer of type v_i with a buyer of type v_{i+1} on this sample path. Given that by assumption it is optimal for the buyer of type v_i to trade in period T, it can only be optimal for a buyer of type v_{i+1} to trade in some period $s \in \{t, \ldots, T\}$.[38] So on the given sample path the expected discounted allocation of the buyer of type v_{i+1} weakly exceeds that of the buyer of type v_i. Summing over all possible sample paths we have $q(v_{i+1}) \geq q(v_i)$. A similar argument applies to sellers. □

Proposition B1. *For any $\boldsymbol{x}$, $V^D(\boldsymbol{x})$ is increasing in v_i for any $i \in \{1, \ldots, n\}$ and decreasing in c_j for any $j \in \{1, \ldots, m\}$. Moreover, letting V^D and $\hat{V}^D$ be the value functions associated with distributions F and $\hat{F}$ respectively, satisfying $\hat{F}(v_i) \leq F(v_i)$ for all i with strict inequality for some, we have $\hat{V}^D \geq V^D$. Likewise, letting V^D and $\overline{V}^D$ be the value functions associated with distributions G and $\overline{G}$ respectively, satisfying $\overline{G}(c_j) \geq G(c_j)$ for all j with strict inequality for some, we have $\overline{V}^D \geq V^D$.*

Proof. We start by noting that to prove each of the desired results for the payoff of the designer, it suffices to prove the results for the value function $V^D(\boldsymbol{x})$ for each $\boldsymbol{x}$. Suppose there is an increase in v_i. Holding the initial optimal policy fixed, this increases the payoff of the designer since any trade involving a buyer of type i now yields a higher reward. Under the new optimal policy the designer's payoff can only increase further, showing that an increase in v_i increases the payoff of the designer. A similar argument applies for a decrease in c_j. Finally, given any sample path, replacing the type of any buyer with some a higher type increases the payoff of the designer. Hence, an increase in F (in the sense of first-order stochastic dominance) leads to an increase in the designer's payoff. A similar argument

[38] Conditional on the buyer of type v_{i+1} optimally remaining until period T on the fixed sample path, if it is optimal to clear a buyer of type v_i in period T it must also be optimal to clear a buyer of type v_{i+1}. However, it may be optimal to clear the buyer of type v_{i+1} in some earlier period.

applies to a decrease (in the sense of first-order stochastic dominance) in G. □

For our main comparative statics results (specifically Corollary 2 and Theorem 3, which compare outcomes under different objectives) to generalize, a new condition, which we call *dynamic regularity* (see (5)), is sufficient whereas Myerson's regularity condition is not. Dynamic regularity ensures that market makers who place a higher value on extracting rent always receive a higher payoff from rematching traders. Thus, if an efficiency-targeting market maker chooses to wait to clear the market in a particular state, so does a profit-maximizing market maker. As in static settings, Myerson's regularity condition is sufficient for pointwise maximization to be implementable in an incentive compatible way. However, in a dynamic setting, it no longer suffices for rent extraction and efficiency to be isomorphic in the sense that the profit-maximizing market maker allocates in the same way as the efficiency-targeting market maker except that his allocation is based on virtual rather than on true types: without dynamic regularity, a profit-maximizing designer may have all sorts of interests in "reshuffling" trades in a dynamic setting. Dynamic regularity guarantees that the isomorphism extends.

Proposition B2. *Under discriminatory market clearing with type spaces that satisfy dynamic regularity, market thickness is increasing in α.*

Proof. Under dynamic regularity we have, for $i \in \{1, \ldots, n-1\}$ and $j \in \{2, \ldots, m\}$,

$$\Phi(v_{i+1}) - \Phi(v_i) > v_{i+1} - v_i \quad \text{and} \quad \Gamma(c_j) - \Gamma(c_{j-1}) > c_j - c_{j-1}$$

which immediately implies that, for $i, i' \in \{1, \ldots, n\}$ with $i > i'$, $j, j' \in \{1, \ldots, m\}$ with $j > j'$ and $\alpha' > \alpha$,

$$\Phi_{\alpha'}(v_i) - \Phi_{\alpha'}(v_{i'}) > \Phi_\alpha(v_i) - \Phi_\alpha(v_{i'}) \quad \text{and} \quad \Gamma_{\alpha'}(c_j) - \Gamma_{\alpha'}(c_{j'}) > \Phi_\alpha(c_j) - \Phi_\alpha(c_{j'}).$$

Thus, if there is an increase in α then it becomes less costly to store trader pairs (since the function Φ_α is decreasing in α and the function Γ_α is increasing in α) and the relative benefits of rematching pairs increases. □

Finally, our results regarding indirect taxation also generalize because under general discrete type spaces, it is still the case that specific taxes distort the relative value of trades while ad valorem taxes do not.

B.2 Uniform market clearing with binary types

Under uniform market clearing, the state space, transition probabilities and reward function of the associated Markov decision process are the same as those of the Markov decision process derived in Section 3.1 for discriminatory market clearing. Uniform market clearing only affects the set of actions available to the designer in a given state. Let $\mathcal{A}'_{\boldsymbol{x}}$ denote the set of actions available to the designer in state $\boldsymbol{x}$ under uniform market clearing. Under discriminatory market clearing we had $\mathcal{A}_{\boldsymbol{x}} = \{(a_E, a_S) : a_E, a_S \in \mathbb{Z}_{\geq 0}, a_S \leq x_E, a_S \leq x_S\}$. However, for the uniform market clearing case the designer can elect only to wait or clear the entire market, implying that $\mathcal{A}'_{\boldsymbol{x}} = \{(x_E, x_S), (0, 0)\}$. Setting $\mathcal{A}' = \cup_{\boldsymbol{x}\in\mathcal{X}}\mathcal{A}'_{\boldsymbol{x}}$, we need to determine the optimal policy of the Markov decision process $\langle \mathcal{X}, \mathcal{A}', P, r, \delta\rangle$. Recall that P specifies the transition probabilities of the Markov chain, $r(\boldsymbol{x})$ specifies the reward earned by the designer when action $\boldsymbol{a} = \boldsymbol{x}$ is implemented and δ is the discount factor. Here, P, r and δ are the same for both discriminatory and uniform market clearing.

In general, we will use the term *threshold policies* to describe any class of policies which can be summarized by one-dimensional sufficient statistics, the thresholds τ. As was the case with discriminatory market clearing, under uniform market clearing we can restrict attention to a class of threshold polices. We then use the structure that threshold policies impose on the market order book to prove that the optimal policy is a threshold policy.

Definition B1. *Given a threshold $\tau \in \mathbb{N}$, the associated* threshold policy π_τ *of the Markov decision process $\langle \mathcal{X}, \mathcal{A}', P, r, \delta\rangle$ is such that*

$$\pi_\tau(\boldsymbol{x}) = \boldsymbol{0} \quad \text{if} \quad r(\boldsymbol{x}) \leq \tau \quad \text{and} \quad \pi_\tau(\boldsymbol{x}) = \boldsymbol{x} \quad \text{if} \quad r(\boldsymbol{x}) > \tau.$$

Under a threshold policy the market maker stores both efficient and suboptimal pairs up to a threshold value of τ. We now describe the associated structure of the order book Markov chain $\{\boldsymbol{Y}_t\}_{t\in\mathbb{N}}$, as illustrated in Figure 7. One can think of the number of stored efficient pairs as the *level* of the Markov chain and the number of stored suboptimal trades as the *phase* of the Markov chain within that level. We include an additional level for the state $\boldsymbol{0}$, denoted by level $\emptyset$. Under the threshold policy τ, $\overline{y}_E = \lfloor \tau \rfloor$ is the maximum number of efficient pairs that can be stored. For $i \in \{0, 1, \ldots, \overline{y}_E\}$, the maximum number of suboptimal pairs stored at level i is $\overline{k}_i = \lfloor (\tau - i)/\Delta_\alpha \rfloor$. Therefore, the order book Markov chain is a level-dependent

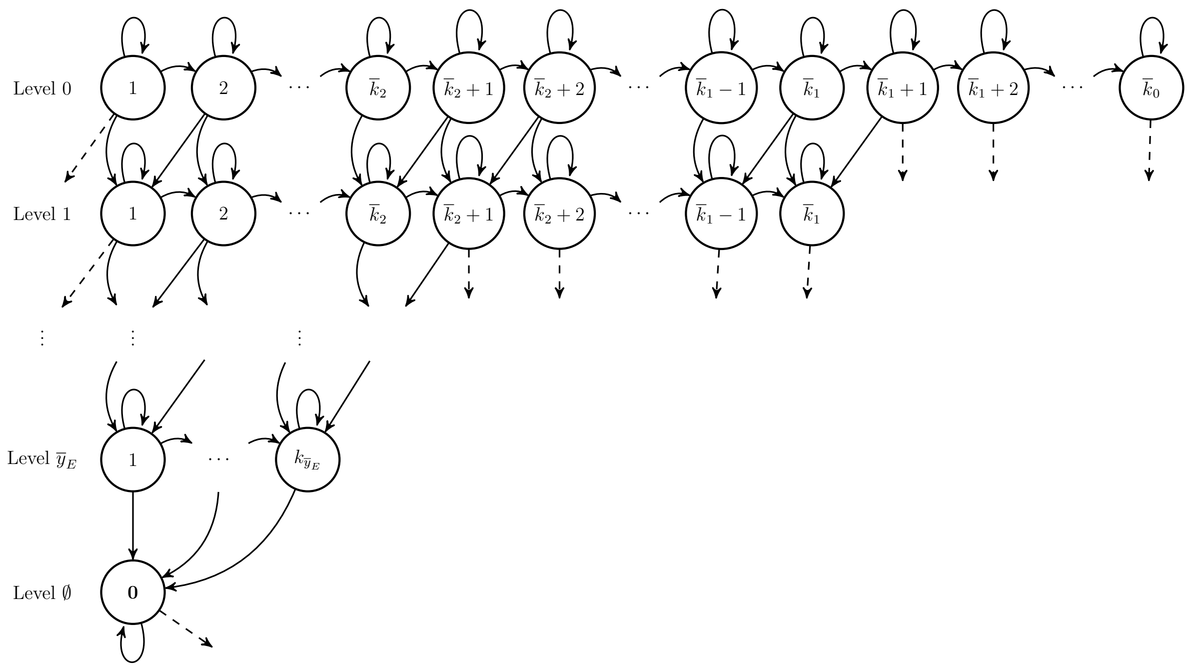

Figure 7: The structure of the quasi-birth-death under the threshold policy with threshold τ. Dashed arrows are used to denote some transitions to and from the state $\mathbf{0}$.

quasi-birth-death process (see, for example, Latouche and Ramaswami (1999)). Similarly to the case of discriminatory market clearing, we can exploit the structure of the order book to show that the optimal market clearing policy is a threshold policy.

Theorem B2. *Under uniform market clearing, the optimal market clearing policy is a threshold policy. It can be implemented using a P-IC and P-IR mechanism.*

The proof of Theorem B2 proceeds in a similar manner to the proof of Theorem 1, using a dynamic programming characterization of the optimal threshold τ^*. Algorithm C2 in Online Appendix C uses the optimal stopping condition derived from the Bellman equation to compute τ^*.

Proof. Since the state space $\mathcal{X}$ is countable, the feasible action sets $\mathcal{A}'_{\boldsymbol{x}}$ are finite for all states $\boldsymbol{x}$ and the reward function is deterministic, a stationary deterministic optimal policy exists (see, for example, Theorem 6.2.10 of Puterman (1994)). Let π^* denote any optimal policy of the Markov decision process $\langle \mathcal{X}, \mathcal{A}', P, r, \delta \rangle$. The optimal policy must clear the market whenever it is in a state of the form $(x_E, 0)$, with $x_E \in \mathbb{N}$. Furthermore, given any fixed number of stored efficient pairs, as the number of stored suboptimal pairs diverges to infinity, the expected time until each additional stored pair is rematched diverges to infinity. Therefore, the benefit of storing each additional suboptimal pair converges to zero, while the immediate reward for clearing a suboptimal pair from the market is fixed at Δ_α. Thus, for a given number of stored efficient pairs, the optimal policy cannot allow an unbounded number of identical suboptimal pairs to accumulate.

It follows that for every $x_E^* \in \mathbb{Z}_{\geq 0}$ there exists a state $\boldsymbol{x}^* = (x_E^*, x_S^*)$ such that $\pi^*(\boldsymbol{x}^*) = \mathbf{0}$ and $\pi^*(x_E^*, x_S^* + 1) = (x_E^*, x_S^* + 1)$. We call such states *cutoff* states. Denote the expected present value of being in the cutoff state $\boldsymbol{x}^*$ under the optimal policy by $V_{\pi^*}^U(\boldsymbol{x}^*)$, the total expected discounted reward earned by the designer in the subsequent period. It is finite because an unbounded number of pairs cannot accumulate under π^* and we are considering a discounted process. For any state $\boldsymbol{x}$, the benefit of waiting to clear the market is increasing in x_S and the benefit of clearing is increasing in $r(\boldsymbol{x})$. Since $r(x_E^* + 1, x_S^*) > r(\boldsymbol{x}^*)$ and $r(x_E^* + 1, x_S^* - 1) > r(\boldsymbol{x}^*)$ it follows that if $\pi^*(x_E^*, x_S^* + 1) = (x_E^*, x_S^* + 1)$, we must also have $\pi^*(x_E^* + 1, x_S^*) = (x_E^* + 1, x_S^*)$ and $\pi^*(x_E^* + 1, x_S^* - 1) = (x_E^* + 1, x_S^* - 1)$. Finally, let $V_{\pi^*}^U(\mathbf{0})$ denote the expected present value of being in the state $\mathbf{0}$ under the optimal policy. The Bellman equation which characterizes $V_{\pi^*}^U(\boldsymbol{x}^*)$ is then given by

$$\begin{aligned} V_{\pi^*}^U(\boldsymbol{x}^*) = \delta \big[& p^2(r(\boldsymbol{x}^*) + 1 + V_{\pi^*}^U(\mathbf{0})) + p(1-p)(r(\boldsymbol{x}^*) + \Delta_\alpha + V_{\pi^*}^U(\mathbf{0})) \\ &+ p(1-p)(r(\boldsymbol{x}^*) + 1 - \Delta_\alpha + V_{\pi^*}^U(\mathbf{0})) + (1-p)^2 V_{\pi^*}^U(\boldsymbol{x}^*)\big]. \end{aligned} \tag{25}$$

If the market is cleared in state $\boldsymbol{x}^*$, the payoff is the immediate reward $r(\boldsymbol{x}^*)$ plus the expected present value of being in the state $\mathbf{0}$. By the principle of the optimality of dynamic programming,

$$V_{\pi^*}^U(\boldsymbol{x}^*) \geq r(\boldsymbol{x}^*) + V_{\pi^*}^U(\mathbf{0}). \tag{26}$$

Notice that the right-hand sides of (25) and (26) depend directly on $\boldsymbol{x}^*$ only through $r(\boldsymbol{x}^*)$. Replace $r(\boldsymbol{x}^*)$ with τ^* in (25) and (26) and suppose (26) holds with equality. Then, for every cutoff state $\boldsymbol{x}^*$, $r(\boldsymbol{x}^*) \leq \tau^*$. Using the definition of τ^*, substituting (26) into (25) and rearranging, it can be shown that τ^* satisfies

$$\tau^* + V_{\pi^*}^U(\mathbf{0}) = \frac{\delta p}{1-\delta}. \tag{27}$$

Thus, for any state $\boldsymbol{x} \in \mathcal{X} \setminus \{(x_E, 0) : x_E \in \mathbb{N}\}$, the market should be cleared if and only if $x_E^* + \Delta_\alpha x_S^* > \tau^*$. Therefore, the optimal policy π^* is a threshold policy, where the threshold $\tau^* \in \mathbb{R}_{\geq 0}$ is characterized by (27).

We now show that the optimal threshold policy can be implemented with a P-IC and P-IR mechanism. Start by constructing a direct allocation rule from the optimal market clearing policy. Let $\hat{h} \in \{\overline{v}, \underline{v}\}^{\mathbb{N}} \times \{\underline{c}, \overline{c}\}^{\mathbb{N}}$ be a realization of the report process and $\hat{h}_t$ denote $\hat{h}$ restricted to its first $2t$ components. Let $\{\tau_j^{\hat{h}}\}_{j \in \mathbb{N}}$ denote the subset of periods such that the designer optimally chooses to clear the market under π^*, given $\hat{h}$ and set $\tau_0^{\hat{h}} = 0$ for convenience. For all $i \in \mathbb{N}$ there exists $j \in \mathbb{N}$ such that $\tau_{j-1}^{\hat{h}} < i \leq \tau_j^{\hat{h}}$. The period $\tau_j^{\hat{h}}$ history of reports can be mapped to $X_{\tau_j^{\hat{h}}}$, the state of $\langle \mathcal{X}, \mathcal{A}', P, r, \delta \rangle$ in period $\tau_j^{\hat{h}}$. Then if buyer i is part of an efficient or a suboptimal pair in period $\tau_j^{\hat{h}}$ we simply set $Q^{B_i}_{\tau_j^{\hat{h}}}(\hat{h}_{\tau_j^{\hat{h}}}) = 1$ and, for all $k \in \mathbb{N} \setminus \{\tau_j^{\hat{h}}\}$, $Q_k^{B_i}(\hat{h}_k) = 0$. Otherwise, we set $Q_k^{B_i}(\hat{h}_k) = 0$ for all $k \in \mathbb{N}$. Proceed analogously for seller i.

Next, we verify the incentive compatibility constraints $q^{B_i}(\overline{v}, \hat{h}_{i-1}) \geq q^{B_i}(\underline{v}, \hat{h}_{i-1})$ and $q^{S_i}(\underline{c}, \hat{h}_{i-1}) \geq q^{S_i}(\overline{c}, \hat{h}_{i-1})$. These constraints hold under π^* since the arrival of a $\overline{v}$ or $\underline{c}$ agent cannot increase the expected number of periods until the next market clearing event (the Markov chain moves to a state with fewer expected transitions between it and the $\mathbf{0}$ state) and $\overline{v}$ and $\underline{c}$ agents are more likely to trade as part of any given market clearing event (these agents have rematching priority over $\underline{v}$ and $\overline{c}$ agents). □

An analogous result to Proposition 3, specifically that the designer's objective function is increasing in p and Δ_0, and τ^* is decreasing in Δ_0, immediately follows from the dynamic programming characterization.[39] Similarly, market thickness, as measured by the optimal threshold τ^*, is increasing in α.

General discrete type spaces Note that we can repeat our analysis in the previous section for uniform market clearing and show that Theorem B1, Proposition B1 and Proposition B2 also hold. The proofs of each of these results for discriminatory market clearing did not rely on the specific structure of the optimal market clearing policy, it exploited general prop-

[39]Notice that τ^* is no longer increasing in $p(1-p)$ since the order book now contains both efficient and suboptimal pairs.

erties of the optimal policies of the Markov decision processes. Therefore, to prove these results it suffices to modify the action space of the general Markov decision process from Section B.1 so that it applies to uniform market clearing.

B.3 Fixed frequency market clearing with binary types

Under fixed frequency market clearing, the state space of the order book Markov chain is given by $\{(y_E, y_S) : 0 \leq y_E + y_S \leq \tau, y_E, y_S \in \mathbb{Z}^2_{\geq 0}\}$. If the market is cleared after τ periods, then both the number of buyers of type $\overline{v}$ present and the number of sellers of type $\overline{c}$ present follow a binomial distribution with τ trials and probability of success p, and likewise for the numbers of buyers of type $\underline{v}$ and sellers of type $\underline{c}$ present. Thus, the market maker's expected discounted payoff is given by

$$V^F_\tau = \frac{\delta^{\tau-1}}{1-\delta^\tau} \sum_{j=0}^{\tau} \sum_{k=0}^{\tau} \binom{\tau}{j} \binom{\tau}{k} \left(\min\{j,k\} + |j-k|\, \Delta_\alpha\right) p^{j+k}(1-p)^{2\tau-j-k}. \tag{28}$$

Here, the market maker can only determine the frequency at which markets are cleared. Thus, the the optimal market clearing policy is trivially a threshold policy, where the market is cleared every τ^* periods. Algorithm C3, which can be found in Online Appendix C, uses this formula to compute the optimal market clearing threshold τ^*.

Once again, we see that the designer's objective function is increasing in p and Δ_0, τ^* is decreasing in Δ_0 and market thickness, as measured by the optimal threshold τ^*, is increasing in α. We also have the following result.

Corollary B1. *Under fixed frequency market clearing, the optimal market clearing policy is a threshold policy. It can be implemented using a P-IC and P-IR mechanism.*

Proof. Under fixed frequency market clearing, threshold policies are trivially optimal. We can repeat the procedure from the proof of Theorem B2 in order to construct a direct allocation rule for fixed frequency market clearing. However, in this case the set of optimal market clearing times is deterministic and given by $\{i\tau^*\}_{i\in\mathbb{N}}$. The constraints $q^{B_i}(\overline{v}) \geq q^{B_i}(\underline{v})$ and $q^{S_i}(\underline{c}) \geq q^{S_i}(\overline{c})$ must then hold since $\overline{v}$ and $\underline{c}$ agents have rematching priority over $\underline{v}$ and $\overline{c}$ agents. □

General discrete type spaces A threshold policy is trivially optimal under fixed-frequency market clearing and the state space of the order book Markov chain is $\mathcal{Y} = \{0, 1, \ldots, \tau^*\}^{n+m}$. Furthermore, just as we had for uniform market clearing, Theorem B1, Proposition B1 and Proposition B2 also hold in this case.

Gains from increased sophistication

We are now in a position to show that the main results from Section 5 (Theorems 2 and 3) generalize to general discrete type spaces.

Theorem B3. *Suppose that $c_1 > v_n$. In the $\delta \to 1$ limit, the relative gains from sophistication are given by*

$$\lim_{\delta \to 1} G_\alpha^{D,U}(\delta) = \lim_{\delta \to 1} G_\alpha^{U,F}(\delta) = 0 < \lim_{\delta \to 1} G_\alpha^{F,0}(\delta).$$

Proof. Let $F^{(-1)}$ and $G^{(-1)}$ denote the respective quantile functions of the distributions F and G. Then in the limit as $\delta \to 1$, per period welfare under discriminatory, uniform and fixed frequency market clearing converges to

$$\int_0^1 \left(F^{(-1)}(1-x) - G^{(-1)}(x)\right) \mathbf{1}(F^{(-1)}(1-x) \geq G^{(-1)}(x))\, dx.$$

Under the assumption that $c_1 > v_n$ per period welfare under instantaneous market clearing is given by the strictly lower value

$$\sum_{i=1}^{n} \sum_{j=1}^{n} (v_i - c_j) \mathbf{1}(v_i \geq c_j) f(v_i) g(c_j)$$

and we immediately have the desired result. □

Theorem B4. *Suppose that the distributions F and G are dynamically regular. Then for all $k \in \{D, U, F\}$ there exists $\delta_k \in [0, 1)$ such that $W^{k,1}(\delta_k) > W^{0,0}(\delta_k)$ for all $\delta > \delta_k$.*

As was done previously, we formally prove the theorem for the $k = D$ case. The other cases are similar.

Proof. First, let $W^\infty(\delta)$ denote welfare under the market clearing policy where all suboptimal trades are stored indefinitely and only efficient trades are executed. Further, let $\mathcal{Y}^{D,0}(\delta)$ and $\mathcal{Y}^{D,1}(\delta)$ denote the set of positive recurrent order book states under welfare-maximizing and profit-maximizing discriminatory market clearing, respectively. Recall that these sets are both increasing in δ. Further, $W^{\infty}(\delta)$ and $W^{0,0}(\delta)$ are continuous, increasing functions of δ and there exists $\tilde{\delta}$ such that $W^{\infty}(\tilde{\delta}) > W^{0,0}(\tilde{\delta})$ because $\mathcal{Y}^{D,0}(\delta) \to \mathbb{Z}_{\geq 0}^{n+m}$ as $\delta \to 1$. By Proposition B2 and our dynamic regularity assumption we have $\mathcal{Y}^{D,0}(\delta) \subseteq \mathcal{Y}^{D,1}(\delta) \subset \mathbb{Z}_{\geq 0}^{n+m}$ for $\delta \in [0,1)$, which in turn implies that $W^{D,1}(\delta) \geq W^{\infty}(\delta)$ for all $\delta \in [0,1)$ (see Lemma A1). It immediately follows that there exists $\delta_D \leq \tilde{\delta}$ such that $W^{D,1}(\delta) > W^{0,0}(\delta)$ for $\delta > \delta_D$. $\square$

C Algorithms

C.1 Discriminatory Market Clearing

We begin by describing an algorithm which may be used to compute the optimal threshold τ^* under discriminatory market clearing. This algorithm exploits the fact that for all $i \in \{0, 1, \ldots, \tau\}$, $V_\tau^D(i)$ is increasing in τ for $\tau \in \{0, 1, \ldots, \tau^* - 1\}$, decreasing in τ for $\tau > \tau^* + 1$ and attains a global maximum at $\tau = \tau^*$ (see Lemma A2).

Algorithm C1. *Begin with the threshold policy characterized by $\tau = 1$ and solve the linear system defined in* (13). *If $V_1^D(1) > \Delta_\alpha + V_1^D(0)$, proceed to step 2. Otherwise, return $\tau^* = 0$. At step i,*

1. *Solve* (13) *with $\tau = i$ to determine $V_i^D(i)$ and $V_i^D(i-1)$.*
2. *If $V_i^D(i) > \Delta_\alpha + V_i^D(i-1)$, proceed to step $i+1$. Otherwise, return $\tau^* = i - 1$.*

Since τ^* is finite, this algorithm must eventually terminate. Algorithm C1 is a simple example of policy iteration. We start with the policy $\tau = 1$ and compute the associated state values. We proceed to iterate over a set of test policies until the optimal policy is reached. With each iteration the test policy is updated based on the optimality condition for the values computed for that test policy. Policy iteration is simple in this case because the set of test policies (which must be the set of all possible optimal policies) has already been refined to the set of threshold policies by Theorem 1.

C.2 Uniform Market Clearing

We next define a similar algorithm that applies to uniform market clearing. However, first we must derive the Bellman equation that characterizes the optimal threshold τ^*. We start by introducing the notation $Z = \{(0,0),(0,1),(1,0),(1,-1)\}$, which captures the set of possible changes to the state $\boldsymbol{y} = (y_E, y_S)$ following the next arrival. Introducing this notation is convenient because it allows us to sum over all possible transitions of the order book Markov chain. Define the function $P_Z : Z \to [0,1]$ by

$$P_Z(1,0) = p^2, \quad P_Z(0,1) = p(1-p), \quad P_Z(1,-1) = p(1-p) \quad \text{and} \quad P_Z(0,0) = (1-p)^2,$$

which gives the probability of each of the changes captured in Z. For example, $(1,-1)$ corresponds to the arrival of a suboptimal pair that results in a stored suboptimal being rematched to create an efficient pair. This occurs with probability $p(1-p)$, provided $y_S > 0$.

Let $V^U_\tau(y_E, y_S)$ denote the expected discounted present value of being in state (y_E, y_S) under the threshold policy with threshold τ. If the state of the market is $(y_E, 0)$ for some $y_E > 0$, the market maker will immediately clear and earn a reward of y_E plus the expected present vale of being in state $\mathbf{0}$. Therefore, we have

$$V^U_\tau(y_E, 0) = y_E + V^U_\tau(\mathbf{0}). \tag{29}$$

Next suppose the market is in any state $\boldsymbol{y} = (y_E, y_S)$ such that $y_S > 0$ and $r(\boldsymbol{y}) < \tau$, where $r(y_E, y_S) = y_E + y_S \Delta_\alpha$ denotes the immediate reward from clearing the market. Under the threshold policy τ, the market maker will earn an immediate reward only when the market reaches a state $\boldsymbol{y}'$ such that $r(\boldsymbol{y}') \geq \tau$. Consequently,

$$\begin{aligned} V^U_\tau(\boldsymbol{y}) = \delta \sum_{\boldsymbol{z} \in Z} P_Z(\boldsymbol{z}) \big[& V^U_\tau(\boldsymbol{y} + \boldsymbol{z}) \mathbb{1}(r(\boldsymbol{y} + \boldsymbol{z}) < \tau) \\ & + (r(\boldsymbol{y} + \boldsymbol{z}) + V^U_\tau(\mathbf{0})) \mathbb{1}(r(\boldsymbol{y} + \boldsymbol{z}) \geq \tau) \big]. \end{aligned} \tag{30}$$

Any threshold policy is characterized by this linear system. As with discriminatory market clearing, this Bellman equation can be used to to derive a stopping condition satisfied by τ^*.

By the proof of Theorem 1, the optimal threshold τ^* is such that for any $x^*_E > 0$ there exists a cutoff state $\boldsymbol{x}^* = (x^*_E, x^*_S)$ with

$$V^U_{\tau^*}(\boldsymbol{x}^*) > r(\boldsymbol{x}^*) + V^U_{\tau^*}(\mathbf{0}) \quad \text{and} \quad V^U_{\tau^*}(x^*_E, x^*_S + 1) \leq r(x^*_E, x^*_S + 1) + V^U_{\tau^*}(\mathbf{0}).$$

That is, a cutoff state is such that the market is optimally cleared if an additional identical suboptimal pair arrives. In the proof of Theorem 2, we show that this implies that the market is then also optimally cleared if an efficient or a non-identical suboptimal pair arrives. Since τ^* applies to all cutoff states, to compute τ^* it suffices to find a single cutoff state. Algorithm C2 determines τ^* by computing the cutoff state $(0, x_S^*)$ using the aforementioned stopping condition.

Algorithm C2. *Begin with the threshold policy characterized by* $\tau = \Delta_\alpha$*, where* Δ_α *is the value of a single suboptimal trade. Solve the linear system defined in* (29) *and* (30)*. If* $V_\tau^U(0,1) \geq \Delta_\alpha + V_\tau^U(\mathbf{0})$*, proceed to step* 2*. Otherwise, return* $\tau^* = 0$*. At step* i*,*

1. *Solve* (29) *and* (30) *with* $\tau = i\Delta_\alpha$ *to determine* $V_\tau^U(0,i)$ *and* $V_\tau^U(\mathbf{0})$*.*

2. *If* $V_\tau^U(0,i) \geq i\Delta_\alpha + V_\tau^U(\mathbf{0})$*, proceed to step* $i+1$*. Otherwise, set* $\tau' = (i-1)\Delta_\alpha$*.*

If $\tau' + \Delta_\alpha < 1$*, return* $\tau^* = \tau'$*. Otherwise, for all* $j \in \mathbb{N}$ *such that* $\tau' + \Delta_\alpha < j$*,*

1. *Set* $k = \lfloor(\tau' + \Delta_\alpha - j)/\Delta_\alpha\rfloor$ *and solve* (29) *and* (30) *with* $\tau = j + k\Delta_\alpha$ *to determine* $V_\tau^U(j,k)$ *and* $V_\tau^U(\mathbf{0})$*.*

2. *If* $V_\tau^U(k,j) \geq j + k\Delta_\alpha + V_\tau^U(\mathbf{0})$ *update* $\tau' = j + k\Delta_\alpha$*.*

Return $\tau^* = \tau'$*.*

Note that depending on the value of δ, Algorithm C2 may not be the most economical algorithm. For example, for larger value of δ, a computationally more efficient algorithm could proceed by initially increasing the candidate threshold by increments of 1 and then increasing the candidate threshold by increments of Δ_α.

C.3 Fixed Frequency Market Clearing

Finally, we have an algorithm to compute τ^* under fixed frequency market clearing.

Algorithm C3. *Begin with the threshold policy characterized by* $\tau = 2$ *and compute* $W^F(2)$ *and* $W^F(1)$ *using* (28)*. If* $W^F(2) \geq W^F(1)$ *proceed to step 2. Otherwise, return* $\tau^* = 1$*. At step* i*,*

1. *Compute* $W^F(i)$ *using* (28).

2. *If* $W^F(i) \geq W^F(i-1)$, *proceed to step* $i+1$. *Otherwise, return* $\tau^* = i-1$.

D Further extensions

We have already considered several extensions of the analyzes in Sections 3, 4 and 5 in this appendix, including general discrete type spaces, uniform market clearing and fixed frequency market clearing. We now briefly discuss several extensions of the arrival process.

D.1 Unpaired arrivals

Under discriminatory market clearing, the assumption that buyers and sellers arrive in pairs can easily be relaxed. To see this, suppose that in every period a buyer of type $\overline{v}$ ($\underline{v}$) arrives with probability $p_1(p_2)$, and with probability $1-p_1-p_2$ no buyer arrives, and likewise for sellers.[40] The designer will optimally store an unbounded number of unpaired efficient types and will store identical suboptimal pairs up to a threshold which can be computed using the methodology described in Section 3.1. Similarly for uniform and fixed frequency market clearing.

D.2 Continuous time

The results of this paper immediately generalize to the case in which buyers and sellers arrive according to a Poisson process. If the intensity of the arrival process is η then the expected inter-arrival time is $1/\eta$ and all of our results generalize if we simply use a discount factor of $\delta^{1/\eta}$. We can also consider the case in which pairs of buyers and sellers arrive according to a more general renewal processes. Let $A(s)$ denote the residual lifetime of the renewal process so that if an arrival occurred at time $t=0$ and a second arrival has not occurred by time $t=s$, then $A(s)$ is the time between s and the next arrival. If the renewal process exhibits the 'new is better than used in expectation" property (that is, if $\boldsymbol{E}[A(s)]$ is decreasing in s; see, for example, Barlow and Proschan (1975)) then our methodology immediately generalizes by simply using a discount factor of $\delta^{\boldsymbol{E}[A(0)]}$. If the renewal process does not have this property,

[40] That is, in every period a seller of type $\underline{c}$ ($\overline{c}$) arrives with probability p_1 (p_2), and with probability $1-p_1-p_2$ no seller arrives.

this is essentially equivalent to considering a time-varying discount factor in our baseline model.

D.3 Group Arrivals

By appropriately updating the transition probabilities of the underlying Markov decision processes, our baseline model could easily accommodate the arrival of groups of agents in each period. This would essentially enable us to consider non-uniform arrival processes which would provide a natural model for markets in which large numbers of buyers and sellers tend to arrive together.[41]

D.4 Multi-Unit Traders

Within our symmetric binary setting it is possible to extend the model to multi-unit traders without sacrificing its amenability to the mechanism design techniques. Specifically, assume that each buyer demands $k \in \mathbb{N}$ units and each seller has the capacity to supply k units. A buyer's type $\theta_{B_t} \in \{0, \ldots, k\}$ is the number of units for which she has a marginal value of $\overline{v}$ while her marginal value for any of the additional units $\max\{k - \theta_{B_t}, 0\}$ is $\underline{v}$. Similarly, seller's type $\theta_{S_t} \in \{0, \ldots, k\}$ is the number of units for which he has a marginal cost of $\underline{c}$ while his marginal cost for producing any of the additional units $\max\{k - \theta_{S_t}, 0\}$ is $\overline{c}$. Assume that buyer types are distributed according to a discrete distribution F with $\text{supp}(f) \subset \{0, \ldots, k\}$ and sellers types are distributed according to some discrete distribution G with $\text{supp}(G) = \{0, \ldots, k\}$.[42] The arrival of the period t buyer and seller is now equivalent to the arrival of $\min\{\theta_{B_t}, \theta_{S_t}\}$ efficient pairs, $|\theta_{B_t} - \theta_{S_t}|$ suboptimal pairs and $k - \max\{\theta_{B_t}, \theta_{S_t}\}$ pairs which cannot trade. Thus, this problem is a special case of the group arrivals extension discussed above.

[41] For example, consider the allocation of university places. Here, students and universities arrive at the market simultaneously prior to the start of the new academic year, essentially creating a static matching problem.

[42] If, for example, buyers valuations for each unit are independent and equal to $\overline{v}$ with probability p and $\underline{v}$ with probability $1 - p$ we obtain a binomial type distribution $\text{Bn}(k, p)$.