Optimal Market Thickness and Clearing*

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

First version: 15 February 2016 This version: February 11, 2018

Abstract

Clearing markets at a lower frequency increases market thickness at the expense of delay. To determine the optimal market clearing policy, we solve a dynamic mechanism design model in which traders arrival stochastically. As the discount factor approaches one, the gains from optimal dynamic market mechanisms relative to instantaneous market clearing go to infinity, while most gains from using a dynamic mechanism are reaped by the simplest dynamic mechanism which clears markets at an optimally chosen frequency. With binary types, efficient incentive compatible trade is possible if storing at least one trade is optimal, in which case the efficient policy can be implemented with budget-balanced posted prices. With a profit maximizing market maker, market thickness is socially excessive, and ad valorem taxes outperform specific taxes.

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 Yeon-Koo Che, Darrell Duffie, Dan Garrett, Scott Kominers, Thomas Mariotti, Leslie Marx, Hayden Melton, Claudio Mezzetti, Paul Milgrom, Carlos Oyarzun, Andy Skrzypacz, Alexander Teytelboym, Leeat Yariv and seminar audiences at the University of Melbourne, the University of Queensland, Ecole Polytechnique Fédérale de Lausanne, 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, and the University of North Carolina at Chapel Hill for comments and suggestions. The second and third author thank the Australian Research Council for funding provided through Laureate Fellowship FL130100039 and the ARC Center of Excellence for Mathematical and Statistical Frontiers (ACEMS). The second author also gratefully acknowledges support provided by the Elizabeth and Vernon Puzey Foundation.

†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 trading associated with continuous trading mechanisms has come under scrutiny on the ground that these trading mechanisms 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 reaction to the arms race and related concerns, a number of major exchanges have recently abandoned the use of continuous-time doubla auction mechanism. For example, the IEX is a stock exchange based in the US that was designed to neutralize the speed advantages required for certain predatory trading strategies by building a 350 microsecond delay into the exchange (Lewis, 2014). Thomson Reuters, one of the two major interbank electronic trading venues in the foreign exchange spot market, abandoned its continuous time market mechanism by introducing a buffer time to deemphasize speed, with the buffer being triggered by trading behavior (Melton, 2017). Similar "speedbump" trading mechanism are currently under consideration at both the New York Stock Exchange (McCrank, 2017b) and the Chicago Stock Exchange (McCrank, 2017a).

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, to clear a market. The fundamental tradeoff is clear: accumulating traders increases market thickness, which is good because it offers more or more valuable opportunities to trade, but bad because it reduces speed and thereby creates costly delay. Traditionally, 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 was designed to execute compatible trades instantaneously, not on the grounds that this would be optimal but because of computational complexity. The basic tradeoff is hard to assess, in parts because the underlying problem is inherently non-stationary, which poses major obstacles for the analysis.

Awareness of the importance of the dynamics of market clearing is also on the rise among for-profit exchanges such as eBay, with recent evidence suggesting that the company foregos substantial profits because it clears instantaneously. Likewise, Uber's reluctance to rely more heavily and systematically on (dynamic) price mechanisms to balance demand and supply may come more under scrutiny as its reliance on behavioral nudges is being criticized.[1] Moreover, as experience with two-sided market design such as the incentive auction accumulates (Milgrom, 2017), these questions are likely to become even more pertinent.[2]

In this paper, we address the problem of determining the optimal market clearing mechanism using a (dynamic) mechanism design approach. In particular, we impose quasilinear payoffs and risk neutrality and assume discrete type spaces, private information about independently distributed private values and costs, stochastic arrival of buyers and sellers, and geometric discounting.

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.[3] Besides its obvious real-world appeal, the setting with private information has the benefit that it

[1]According to a recent New York Times article (How Uber Uses Psychological Tricks to Push Its Drivers' Buttons), Uber's vice president of product is reluctant to increase consumer prices (use "surge" prices): "For us, it's better not to surge. ... If we don't surge, we can produce more rides."

[2]For a recent overview of the literature on two-sided market design and the questions arising there, see also Loertscher et al. (2015).

[3]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).

neither presumes nor precludes efficiency.[4] It also makes the problem of revenue generation interesting in a plausible way. We depart from the standard Myersonian setup by assuming discrete types – indeed, in the baseline model we assume binary types and pairwise arrival of buyers and sellers – to make the state space and the model tractable. We will discuss how discreteness affects the results as we go. Lastly, 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.

In our mechanism design analysis, we first impose appropriate notions of incentive compatibility and individual rationality and derive expected revenue. Then we relax the global incentive compatibility constraints and determine the allocation rule that maximizes the designer's objective function (as implied by the local incentive compatibility and individual rationality constraints). Finally, we verify that this pointwise maximizer satisfies the global incentive compatibility constraints. In contrast to standard, static mechanism design settings, where the pointwise maximizer is often trivial, considerable work goes into the derivation of the allocation rule. We solve this problem by mapping the allocation problem to a Markov decision process, which allows us to account for the endogeneity of future states and paths. We show that a particularly simple type of policy, a so-called *threshold* policy, is optimal.[5] That this allocation rule indeed permits incentive compatible transfers then follows from its natural and intuitive monotonicity properties.

A key result from our analysis is that the impossibility of efficient, incentive compatible and individually rational trade with privately informed agents is overcome 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 Coase (1960), Vickrey (1961), Hurwicz (1972), and Myerson and Satterthwaite (1983). Because of its simplicity, implementation via posted prices is also

[4]That private information is an at times 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).

[5]When type spaces are richer, threshold policies are used to construct a finite partition of the state space. This allows the optimal market clearing policy to be determined using a simple implementation of the standard policy iteration algorithm (see Howard, 1960)).

appealing for practical purposes.[6] While our results have a Coasian flavour in that they show that inefficient initial allocations can be resolved efficiently 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 initial allocation cannot always be resolved efficiently.[7] Our analysis shows that the transactions costs due to private information and stochastic arrival of traders can be overcome through an appropriately designed, centralized market mechanism.[8]

We first derive the optimal mechanism and market clearing policy when different traders present in the market at the same time may be cleared at different times, which we refer to as *discriminatory* market clearing. In many real-world market places, so much flexibility is not deemed possible. 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, 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. We derive the optimal market mechanism under all three forms of market clearing and contrast the outcomes these imply with those under instantaneous clearing (that is, when compatible orders that arrive at the same time are cleared instantaneously and no other trades occur). We show that asymptotically, as the discount factor approaches one, discounted expected welfare gains relative to those under

[6] 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.

[7] 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.

[8] 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.

instantaneous trade under all three forms of dynamic market clearing (i.e. discriminatory, uniform, fixed frequency) converge and are unbounded. This means that, asymptotically, the exact nature of dynamic market clearing is of second order importance and that the welfare gains from dynamic market clearing can be substantial, indicating potentially new and important scope for market design in dynamic settings.

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 period ex post efficient exchange that never stores any trades.

Moreover, we show that an ad valorem tax imposed on the profit of 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, for example, fared prominently in policy debates following the Global Financial Crisis.

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), consider 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 afore-mentioned papers explicitly address the optimal timing problem[9] nor do they consider varying degrees of market sophistication or compare welfare and profit maximization.[10]

Methodologically, our paper is closely related to the literature on dynamic matching without monetary transfers. In recent work that builds 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. The efficient algorithm Baccara et al. (2016) determine for a dynamic matching model is similar to the pointwise maximizer that is optimal for our designer. However, besides not allowing for transfers, their model does not have geometric discount nor quasilinear utility. Consequently, it does not admit the kind of analyses that are central to our paper, such as the possibility or impossibility of efficient trade, optimal rent extraction, comparisons of the performance of a dynamic for-profit exchange with that of a myopic, periodically ex post efficient exchange, posted prices, or indirect taxes.

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 have to center around the question of transaction costs. Vickrey (1961), Hurwicz (1972) and Myerson and Satterthwaite (1983) argued forcefully that private 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. Our paper contributes to this debate, which thus far was pursued largely assuming static settings, 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.

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

[10]There is a vast literature on intermediation in financial markets; see, for example, Mendelson (1982) or Kelly and Yudovina (2016) and references therein. Considering a sub-optimal market mechanism, Malamud and Rostek (2016) show that decentralized exchanges may outperform centralized ones. In contrast to our paper, this literature does not consider agents' incentives in two-sided intermediated markets. It also leaves aside the possibility that the market maker may be interested in generating profits.

The remainder of this paper is organized as follows. Section 2 introduces the setup. In Section 3, we solve the mechanism design problem, discuss key properties of the optimal mechanism and relate the dynamically efficient allocation rule to price posting and the (im)possibility of efficient trade. Section 4 analyzes compares the performance of the optimal dynamic mechanism to two less sophisticated dynamic mechanisms and instantaneous market clearing, while Section 5 contains extensions. Section 5 concludes. All proofs and algorithms are provided in the Appendix.

2 Setup

In this section, we lay out the baseline model. 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 the designer faces.

2.1 Agents, Types, and the Arrival Process

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. 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 assume that agents can only trade via the designer's platform and the value of agents' outside option of not participating is zero.

We also assume that buyers are of type $\overline{v}$ or $\underline{v}$ with probability p or $1 - p$, respectively, satisfying $\overline{v} > \underline{v}$, and sellers are of type $\underline{c}$ with probability p and $\overline{c}$ with probability $1 - p$, respectively, with $\underline{c} < \overline{c}$. Throughout this paper, 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*. We make the standard mechanism design assumption that agent types are private information while the parameters $\underline{v}$, $\overline{v}$, $\underline{c}$, $\overline{c}$ and p are common knowledge. Information is elicited from period t agents once, when they report their types upon arrival. For our setup with persistent private values, this is without loss of generality. 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

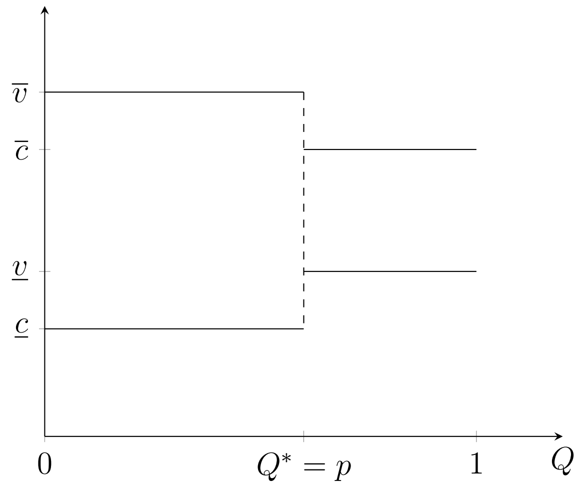

Figure 1: In every period, a buyer-seller pair arrives. Buyers and sellers draw their values and costs independently from the distributions $\{\underline{v}, \overline{v}\}$ and $\{\underline{c}, \overline{c}\}$, respectively, with $\underline{c} < \underline{v} < \overline{c} < \overline{v}$, probability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

type $\overline{v}$ and buyers of type $\underline{c}$. The Walrasian equilibrium for this setup is illustrated in Figure 1 under the assumption

$$\overline{v} > \overline{c} > \underline{v} > \underline{c}, \tag{1}$$

which we maintain throughout the paper. In fact, with the exception of Subsection 5.2, which generalizes the type space, we normalize $\overline{v} = 1$ and $\underline{c} = 0$ and impose additional symmetry by setting

$$\underline{v} = \Delta_0 \quad \text{and} \quad \overline{c} = 1 - \Delta_0, \tag{2}$$

where $\Delta_0 \in (0, 1/2)$ because of (1) and the normalization $\overline{v} = 1$ and $\underline{c} = 0$. Assumption (1) makes sure that in the static setup with a continuum of agents high-value buyers and low-cost sellers trade in the Walrasian market and low-value buyers and high-cost sellers remain inactive. It also implies 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, which in our dynamic setting upon the arrival of a pair $(\overline{v}, \overline{c})$ (or $(\underline{v}, \underline{c})$) 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. Imposing the symmetry condition (2) then allows us to focus on this tradeoff without being distracted by additional complications due to asymmetries.

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 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 some revenue constraint $R \geq \underline{R}$,[11] 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

$$\alpha R(\langle \boldsymbol{Q}, \boldsymbol{M} \rangle) + (1 - \alpha) W(\langle \boldsymbol{Q}, \boldsymbol{M} \rangle), \tag{3}$$

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

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 a period t agent that reports $\hat{\theta}$. Similarly, let $m(\hat{\theta})$ denote the expected payment made (received) by a period t buyer (seller) that reports $\hat{\theta}$. Then the *Bayesian incentive compatibility (BIC)* and *individual rationality (IR)* constraints require that for all $v \in \{\underline{v}, \overline{v}\}$ and $c \in \{\underline{c}, \overline{c}\}$,

$$v = \underset{\hat{\theta} \in \{\underline{v}, \overline{v}\}}{\arg\max} \left\{ v q(\hat{\theta}) - m(\hat{\theta}) \right\} \quad \text{and} \quad c = \underset{\hat{\theta} \in \{\underline{c}, \overline{c}\}}{\arg\max} \left\{ m(\hat{\theta}) - c q(\hat{\theta}) \right\}, \tag{BIC}$$

[11] Or, equivalently, maximize R subject to some welfare constraint $W \geq \underline{W}$.

and

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

For any $\alpha > 0$, it is well known that individual rationality constraints bind for the worst-off types (i.e. buyers of type $\underline{v}$ and sellers of type $\overline{c}$) and that, with binary types, the incentive compatibility constraints bind for the efficient ("best-off") types (i.e. $\overline{v}$ and $\underline{c}$).[12,13] Thus, we have $m(\underline{v}) = \underline{v}q(\underline{v})$, $m(\overline{c}) = \overline{c}q(\overline{c})$, $\overline{v}q(\overline{v}) - m(\overline{v}) = \overline{v}q(\underline{v}) - m(\underline{v})$ and $m(\underline{c}) - \underline{c}q(\underline{c}) = m(\overline{c}) - \underline{c}q(\overline{c})$, giving

$$m(\overline{v}) = \overline{v}(q(\overline{v}) - q(\underline{v})) + \underline{v}q(\underline{v}) \quad \text{and} \quad m(\underline{c}) = \underline{c}(q(\underline{c}) - q(\overline{c})) + \overline{c}q(\overline{c}). \tag{4}$$

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

Below we will show that, in our setting, alternative, stronger notions of incentive compatibility that have been discussed in the literature are equivalent to (BIC). The *interim incentive compatibility constraints (i-IC)* (see e.g. Bergemann and Välimäki, 2010) require that truthful reporting is optimal for every period t agent and every history $\boldsymbol{h}_{t-1} \in \mathcal{H}_{t-1} := (\{\overline{v}, \underline{v}\} \times \{\underline{c}, \overline{c}\})^{t-1}$, 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 $\boldsymbol{h}_{t-1} \in \mathcal{H}_{t-1}$, $v \in \{\underline{v}, \overline{v}\}$ and $c \in \{\underline{c}, \overline{c}\}$, (i-IC) then requires

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

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

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

[12] 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. We avoid this indeterminacy be treating $\alpha = 0$ as the limit of $\alpha \to 0$, which implies that for a given allocation rule the incentive-compatible transfers are revenue maximizing, which in turn implies that the individual rationality constraints binds.

[13] More generally, with discrete type spaces, the individual rationality constraints bind for worst-off types. For all other types, the incentive compatibility constraints bind locally downward for buyers and bind locally upward for sellers; see Elkind (2007).

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 $\boldsymbol{h}_{t-1} \in \mathcal{H}_{t-1}$, $v \in \{\underline{v}, \overline{v}\}$ and $c \in \{\underline{c}, \overline{c}\}$, (P-IC) requires

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

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

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

Without loss of generality, we restrict ourselves to direct, truthful and deterministic mechanisms.[14] 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 $\boldsymbol{M} = \{\boldsymbol{M}_t\}_{t\in\mathbb{N}}$. The period t allocation rule $\boldsymbol{Q}_t : \mathcal{H}^t \to \{0, 1\}^{2t}$ maps the period t history of agents reports $\boldsymbol{h}_t$ to the set of period t allocations, where $\mathcal{H}_t := (\{\overline{v}, \underline{v}\} \times \{\underline{c}, \overline{c}\})^t$ with typical element $\boldsymbol{h}_t = ((v_1, c_1), \ldots, (v_t, c_t))$. Similarly, we have a period t transfer rule $\boldsymbol{M}_t : \mathcal{H}_t \to \mathbb{R}^{2t}$.

When the period t report history is given by $\boldsymbol{h}_t$, the respective period t allocations of buyer and seller $i \in \{1, \ldots, t\}$ are $Q_t^{B_i}(\boldsymbol{h}_t)$ 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 made by B_i and S_i in period t given $\boldsymbol{h}_t$. Feasibility requires that, for all $t \in \mathbb{N}$ and all $\boldsymbol{h}_t \in \mathcal{H}_t$,

$$\sum_{i=1}^{t} Q_t^{B_i}(\boldsymbol{h}_t) \leq \sum_{i=1}^{t} Q_t^{S_i}(\boldsymbol{h}_t) \tag{5}$$

and, for all $i \in \{1, \ldots, t\}$, $\sum_{j=i}^{t} Q_j^{B_i}(\boldsymbol{h}_j) \leq 1$ and $\sum_{j=i}^{t} Q_j^{S_i}(\boldsymbol{h}_j) \leq 1$. Of course, (5) will hold with equality under an optimal mechanism.

[14] For direct, truthful mechanisms, no generality is lost because of the revelation principle. The restriction to deterministic mechanisms is without loss of generality because we will find that optimal mechanisms are deterministic up to the resolution of ties, which can be broken using deterministic rules.

3 Optimal Mechanisms

We first derive the allocation rule that maximizes the designer's objective function as implied by incentive compatibility and individual rationality pointwise. Then we verify that the pointwise maximizer permits incentive compatibility. 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 Virtual Types

Denote by $\Phi(\overline{v}) := \overline{v}$ and $\Gamma(\underline{c}) := \underline{c}$ the virtual value and virtual cost, respectively, of the efficient types and by

$$\Phi(\underline{v}) := \underline{v} - \frac{p}{1-p}(\overline{v} - \underline{v}) \quad \text{and} \quad \Gamma(\overline{c}) := \overline{c} + \frac{p}{1-p}(\overline{c} - \underline{c}) \tag{6}$$

the virtual value and virtual cost, respectively, of the inefficient types. 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).

Expected discounted social welfare under any direct, truthful mechanism that implements the allocation rule $\boldsymbol{Q}$ is given by

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

Using standard mechanism design arguments, the designer's expected discounted profit under 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 $\boldsymbol{Q}$ satisfying (BIC) and (IR) is given by*

$$R(\boldsymbol{Q}) = \sum_{t=1}^{\infty} \sum_{i=1}^{t} \sum_{\boldsymbol{h}_t \in \mathcal{H}_t} \delta^{t-1} \left(\Phi(v_i) Q_t^{B_i}(\boldsymbol{h}_t) - \Gamma(c_i) Q_t^{S_i}(\boldsymbol{h}_t) \right) \mathbb{P}(\boldsymbol{H}_t = \boldsymbol{h}_t). \tag{8}$$

Furthermore, $R(\boldsymbol{Q})$ given in (8) is also the expected discounted profit under any direct mechanism with allocation rule $\boldsymbol{Q}$ satisfying (i-IC) and (i-IR) respectively (P-IC) and (P-IR).

Using (7) and (8), we can now rewrite the Ramsey objective (3), incorporating incentive compatibility and individual rationality constraints, as

$$\alpha R(\boldsymbol{Q}) + (1-\alpha)W(\boldsymbol{Q}) = \sum_{t=1}^{\infty}\sum_{i=1}^{t}\sum_{\boldsymbol{h}_t\in\mathcal{H}_t} \delta^{t-1}\left(\Phi_\alpha(v_i)Q_t^{B_i}(\boldsymbol{h}_t) - \Gamma_\alpha(c_i)Q_t^{S_i}(\boldsymbol{h}_t)\right)\mathbb{P}(\boldsymbol{H}_t = \boldsymbol{h}_t), \tag{9}$$

where, for $\alpha \in [0,1]$, $v \in \{\underline{v}, \overline{v}\}$ and $c \in \{\underline{c}, \overline{c}\}$,

$$\Phi_\alpha(v) \coloneqq (1-\alpha)v + \alpha\Phi(v) \quad \text{and} \quad \Gamma_\alpha(c) \coloneqq (1-\alpha)c + \alpha\Gamma(c) \tag{10}$$

are the weighted virtual types. The designer's problem is now to determine the allocation rule $\boldsymbol{Q}_\alpha$ that maximizes (9), subject to the appropriate incentive constraints.

In this section, we allow the designer to make market clearing contingent on the state of the market, and to partially clear the market (that is, the clear some trades at any point in time while storing others). We refer to this as *discriminatory market clearing*.

3.2 Deriving the Optimal Allocation Rule

We now derive the allocation rule $\boldsymbol{Q}_\alpha$ that maximizes (9), temporarily neglecting the constraint that this rule be incentive compatible. We will later verify that the relevant incentive constraints hold. The relaxed optimization problem can be rewritten in terms of a Markov decision process.

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

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

where the inequality is strict for $\alpha > 0$. Notice also that the designer only wants to induce trade for buyers with high valuations and sellers with low costs if $\Delta_\alpha < 0$ holds. Thus, if $\Delta_\alpha < 0$, the derivation of the optimal allocation rule is trivial.[15] For the remainder of this subsection we assume that the parameters α, p and Δ_0 are such that

$$\Delta_\alpha > 0. \tag{11}$$

[15]It 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.

Since $\Delta_\alpha \le \Delta_0 < 1/2$, 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 $2\Delta_\alpha$.

Under discriminatory market clearing, the designer's problem is to determine which pairs should be cleared from the market in each period. When a pair which 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 call a pair $(\overline{v}, \underline{c})$ pair *efficient*, the pairs $(\overline{v}, \overline{c})$ and $(\underline{v}, \underline{c})$ *suboptimal* and a $(\underline{v}, \overline{c})$ pair *inefficient*. The underlying state at date 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. Notice 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 binary and given by $\mathcal{X} \coloneqq \{(x_E, x_S) : x_E, x_S \in \mathbb{Z}_{\ge 0}\}$, 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.

Denote by $\mathcal{A}_{\boldsymbol{x}}$ the set of *actions* available to the designer in state $\boldsymbol{x}$ and let $\mathcal{A} = \cup_{\boldsymbol{x} \in \mathcal{X}} \mathcal{A}_{\boldsymbol{x}}$. Under discriminatory market clearing we have $\mathcal{A}_{\boldsymbol{x}} = \{(a_E, a_S) : a_E, a_S \in \mathbb{Z}_{\ge 0}, a_E \le x_E, a_S \le x_S\}$, 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

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

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 (9). Thus, the designer's relaxed optimization problem reduces to determining the *optimal policy* π^* of $\langle \mathcal{X}, \mathcal{A}, P, r, \delta \rangle$.

3.3 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[16] and that

[16]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.

the arrival process is stationary. Threshold policies are analogous to policies induced by a Gittins (1979) index, that apply to multi-armed bandit problems. We also prove that the optimal threshold policy, which solves the designer's relaxed optimization problem, also satisfies the global incentive compatibility constraints.

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

We now briefly discuss the stationarity and uniqueness of the optimal allocation rule and mechanism. Observe that the optimal market clearing policy is unique up to the identities of agents that are cleared from the market when more than one of 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 the designer's payoff is up invariant to the treatment of individual agents that arrive as part of a suboptimal pair. 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.[17] However, sationariy of the optimal market clearing policy does not imply that the optimal allocation rule is stationary. In fact, the optimal market clearing policy could be implemented with a non-stationary queueing protocol over agents of the same type. Nevertheless, the optimal policy can always be implemented with a stationary queueing protocol, which is an attractive feature if it is deemed undesirable that the treatment of traders' depends on the period in which they arrive.

For $\alpha = 0$, we refer to optimal policy of Theorem 1 as the *first-best* policy because it does not impose constraints beyond those inherent in the problem. Although Theorem 1 does not immediately allow us to identify the optimal market clearing policy, 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 next to characterize the optimal threshold policy τ^*. In particular, each threshold policy 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

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

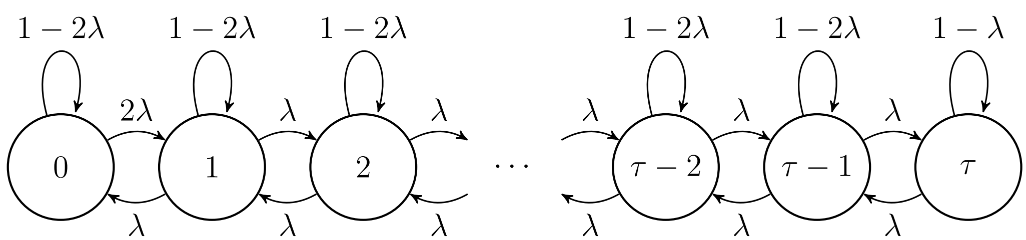

Figure 2: The Markov chain over the number of stored suboptimal pairs induced by the optimal policy π^* under discriminatory market clearing with $\lambda = p(1-p)$.

distribution of this Markov chain is straightforward.

Proposition 2. *Under the discriminatory market clearing threshold policy with threshold τ, the stationary distribution $\boldsymbol{\kappa}$ of the Markov chain $\{Y_t\}_{t\in\mathbb{N}}$ is given by*

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

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

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

By determining the immediate reward earned by the designer transitions take place in the order book Markov chain, we can write the Bellman equation associated with the Markov decision process. Take any $y \in \{0, 1, \dots, \tau\}$ 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 $y \in \{1, \dots, \tau-1\}$, characterized by the Bellman equation

$$V_\tau^D(y) = \delta\left[p^2(1 + V_\tau^D(y)) + p(1-p)(1 + V_\tau^D(y-1) + V_\tau^D(y+1)) + (1-p)^2 V_\tau^D(y)\right], \quad (12)$$

with boundary conditions

$$V_\tau^D(0) = \delta\left[p^2(1 + V_\tau^D(0)) + 2p(1-p)V_\tau^D(1) + (1-p)^2 V_\tau^D(0)\right] \quad (13)$$

and

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

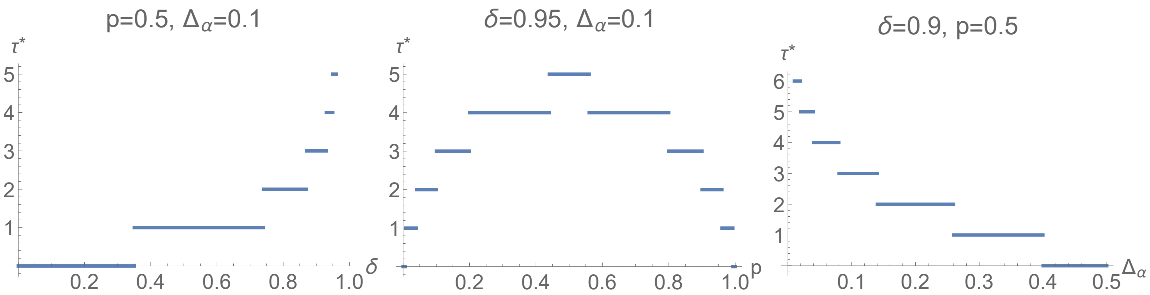

(a) Increasing in δ (b) Increasing in $p(1-p)$ (c) Decreasing in Δ_α

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

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

$$V^D_{\tau^*}(\tau^*) > \Delta_\alpha + V^D_{\tau^*}(\tau^* - 1) \quad \text{and} \quad V^D_{\tau^*+1}(\tau^* + 1) \leq \Delta_\alpha + V^D_{\tau^*+1}(\tau^*). \tag{14}$$

To compute the optimal threshold, one can start with the threshold policy given by $\tau = 1$, check condition (14) and iterate. Algorithm 1 in Appendix B formalizes this procedure.

Proposition 3. *Under the optimal policy, the designer's objective function is increasing in δ, p and Δ_α. The optimal threshold τ^* is increasing in δ and in $p(1-p)$ and decreasing in Δ_α.*

Proposition 3 says that the designer's total expected discounted payoff is increasing in the discount factor, the value of suboptimal trades and the probability of efficient types arriving on each side of the market. The comparative statics of τ^* with respect to δ, $p(1-p)$ and Δ_α are illustrated in Figure 3. Intuitively, τ^* increases as the cost of storing traders decreases. Hence, it is increasing in δ and decreasing in Δ_α. 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. In a given period, this probability is $p(1-p)$. Of course, τ^* is a straightforward measure of market thickness. Interestingly, Proposition 3 has the following corollary.

Corollary 1. *Market thickness as measured by τ^* is increasing in α.*

Corollary 1 is reminiscent of Hotelling's (1931) finding that a monopolist extracts an exhaustible resource at a slower rate than a perfectly competitive industry.[18] As is the case

[18]Corollary 1 does not necessarily extend to finite horizon models with richer type spaces. For example,

in static environments, these distortions arise under the optimal mechanism as a means of reducing the informational rents of agents.[19]

Dynamic Efficiency Corollary 1 has important implications. In the perfectly patient limit (i.e. as $\delta \to 1$), which is equivalent to a static setup with a continuum of traders, a trade is executed if and only if it is efficient.[20] 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.[21]

However, the efficient outcome is different when $\delta < 1$. Under the first-best policy, 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, what is inefficient in a static setting is an integral part of efficiency in a dynamic setting. Moreover, for a fixed discount factor, by Corollary 1 the fewer are such apparently inefficient suboptimal trades, the greater is the market maker's rent extraction.

Another fundamental difference to static setups is that in the dynamic setting efficiency is not a distribution-free concept because the optimal mechanism depends on p. This has implications for implementation in "detail-free" environments as discussed in Section 3.5.

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 his 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)

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

[20]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.

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

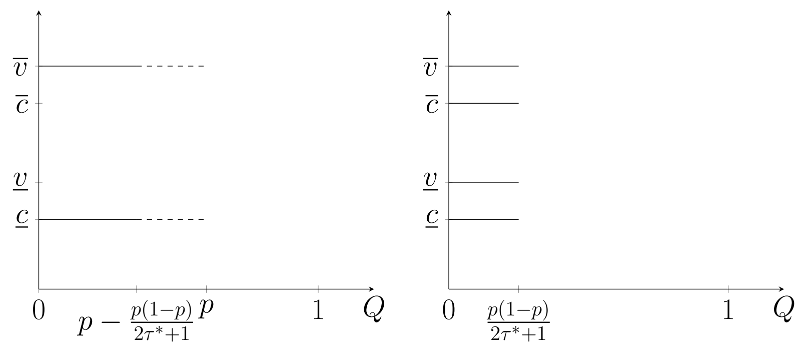

(a) Efficient trade (b) Suboptimal trades

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

3.4 Posted Prices and the Possibility of Efficient Trade

The optimal mechanism asks agents to directly report types and makes payments and allocations 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 now look at posted prices mechanisms and investigate their relationship with the first-best policy and first-best mechanism when $\alpha = 0$.

We begin with the following simple *posted price mechanism*, under which trade occurs at a uniform price that is posted at the start of each period. We will show that an appropriately structured posted price mechanism implements the first-best policy whenever it is optimal to store at least one suboptimal pair under first-best. Thereby, our paper relates dynamic market making both to the classic literature on the (im)possibility of efficient trade and to price posting. Furthermore, using the posted price mechanism allows us to characterize the stationary distribution of prices when price posting implements the optimal allocation rule and to provide a measure of an agent's price impact.[22] As $\delta \to 1$, we show that the variance of the price distribution vanishes and that the price distribution converges to a degenerate distribution with probability mass 1 on a Walrasian price in the static model with a continuum of traders. We conclude this subsection by discussing the relation of this convergence result to the results in the literature on the microfoundation of Walrasian equilibrium based on search and random matching models.

[22]Equivalently, this be interpreted as a measure of the depth of a market.

Definition 2. *A* posted price mechanism *proceeds as follows. At the start of each period t the designer posts, as a function of the state of the order book, a price p_B for buyers and a price p_S for sellers. The period t agents then arrive and all agents observe the order book and the posted prices before making a report of $\rho \in \{0, 1, [0, 1]\}$, where 1 indicates that the agent accepts the posted price, 0 indicates that the agent rejects the posted price and report of $[0, 1]$ 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. Thus, a posted price mechanism is characterized by the pricing rules p_B and p_S and the queueing protocol.*

To analyze the incentive properties of posted price mechanisms, we make the simplifying assumption that the designer removes any agent that makes inconsistent reports from the market.

A Budget Balanced Posted Price Mechanism

Consider the following posted price mechanism, which we refer to as the *balanced budget* posted price mechanism with threshold $\tau \geq 1$.

Definition 3. *Let a threshold $\tau \geq 1$ be given. We define the associated* balanced budget posted price mechanism *as follows. Suppose that the number of suboptimal pairs stored in the order book is $y < \tau$. Then designer posts prices $p_B = p_S = 1/2$. Similarly, suppose that the number of suboptimal pairs stored in the order book is $y = \tau$. Then the designer posts prices $p_B = p_S = \Delta_0$ if $(\underline{v}, \underline{c})$ pairs are stored and $p_B = p_S = 1 - \Delta_0$ if $(\overline{v}, \overline{c})$ pairs are stored. We assume that a last-come-first-served queueing protocol is used to determine the order in which agents are cleared from the market. We further break market-clearing ties by assuming that if arriving agents are indifferent between accepting and rejecting the posted prices, they are cleared from the market, while indifferent pairs of stored agents are not cleared from the market.*

By construction, the balanced budget posted price mechanism does not run a deficit. Under truthful reporting the mechanism immediately executes efficient trades and does not execute any suboptimal trades when less than τ identical suboptimal pairs are stored. Once

τ $(\overline{v}, \overline{c})$ pairs are stored, the design 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 $(\underline{v}, \underline{c})$ trades 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 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 and $\alpha = 0$, the first-best policy executes all trades that generate positive surplus. However, the downside to the mechanism that implements the first-best policy when $\alpha = 0$ and $\delta = 0$ is that, depending on the parametrization, it may run a deficit.[23] Interestingly, it turns out that there is a tight connection between the budget balanced posted price mechanism and the first-best 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 we have already discussed, the budget balanced posted price mechanism implements the first-best allocation rule, provided agents report truthfully. So we only need to check the relevant 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 be cleared from 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,

[23]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, thus $p(2\Delta_0 - p)$, which is negative for $p > 2\Delta_0$.

the incentive constraints are satisfied for buyers of type $\overline{v}$. Since a similar argument applies to sellers of type $\underline{c}$, we are done. Recall that in our discussion of the first-best 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.

Although the budget balanced posted price mechanism in not the first-best mechanism, because the first-best mechanism must generate weakly more profit for the market marker, we immediately have the following corollary to Proposition 4.

Corollary 2. *The first-best mechanism (i.e. the optimal mechanism with* $\alpha = 0$*) does not run a deficit if* $\tau^* > 0$*.*

As noted in footnote 23, 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.[24] 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 markets that are constrained not to run a deficit grow large. 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 can occur for δ as small as 0.36 as illustrated in Figure 3.[25] 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.[26]

[24]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.

[25]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 2 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 23), is also sufficient for a deficit with $N = 2$ pairs present.

[26]As noted in the introduction (see in particular footnote 6), 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

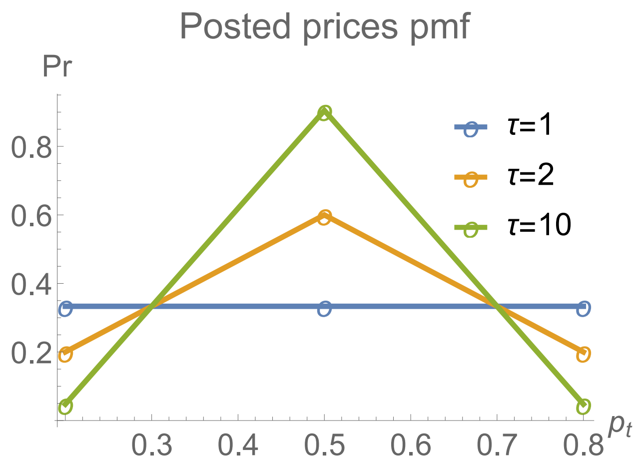

Figure 5: The probability mass function of the posted price under the budget balanced posted price mechanism for $\Delta_0 = 0.2$ implying $p_t \in \{0.2, 0.5, 0.8\}$.

A similar result holds for more general discrete type spaces (see Proposition 6 in Section 5). However, as the type spaces become increasingly rich, a larger value of δ is required for the efficient dynamic mechanism to create market thickness sufficient for the deficit to vanish. And of course, by the analysis of Gresik and Satterthwaite (1989), if we have continuous type spaces the impossibility result vanishes only in the limit as $\delta \to 1$.

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 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

$$\mathbb{P}(p_t = \underline{v}) = \mathbb{P}(p_t = \overline{c}) = \frac{1}{2\tau^* + 1} \quad \text{and} \quad \mathbb{P}(p_t = 1/2) = \frac{2\tau^* - 1}{2\tau^* + 1}, \tag{15}$$

where the equalities in (15) follow from the stationary distribution given in Proposition 2. An illustration of this distribution is given in Figure 5. It further follows that the stationary variance of the posted prices is

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

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).

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

Corollary 3. $\mathbb{P}(p_t = 1/2)$ *increases in* δ *and* $p(1-p)$ *and decreases in* Δ_0. $\mathrm{Var}(p_t)$ *decreases in* δ *and* $p(1-p)$. *Moreover,*

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

The first part of the corollary says that the distribution shifts ever more weight to the static Walrasian price as the discount factor increases and the probability of a contemporaneous mismatch (i.e. $p(1-p)$) 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 extreme, thereby all else equal increasing the variance, but on the other hand decrease the difference between the lowest and the highest price in the support, i.e. between Δ_0 and $1 - \Delta_0$.

The limit results in the second part of Corollary 3 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 parameterized 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 (that is, $\overline{c}$ if he is a buyer, $\underline{v}$ if he is 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 he is part of another suboptimal pair. Therefore, defined in this way, an agent's likely price impact is

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

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.

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 post-price mechanism that implements the efficient allocation, whilst maximizing profit for the market maker.

Definition 4. *We define the associated* optimal posted price mechanism *with threshold* $\tau \geq 1$ *as follows. Suppose that the number of suboptimal pairs stored in the order book is* $y < \tau$. *Then designer posts prices* $p_B = 1$ *and* $p_S = 0$. *Similarly, suppose that the number of suboptimal pairs stored in the order book is* $y = \tau$. *Then the designer posts prices* $p_B = \Delta_0$ *and* $p_S = 0$ *if* $(\underline{v}, \underline{c})$ *pairs are stored and* $p_B = 1$ *and* $p_S = 1 - \Delta_0$ *if* $(\overline{v}, \overline{c})$ *pairs are stored. We again assume that a last-come-first-served queueing protocol is used to determine the order in which agents are cleared from the market. We further break market-clearing ties by assuming that if arriving agents are indifferent between accepting and rejecting the posted prices, they are cleared from the market, while indifferent pairs of stored agents are not cleared from the market.*

Proposition 5. *Within the class of posted price mechanisms identified in Definition 2, the optimal posted price mechanism with* $\tau = \tau^*$ *provides a P-IC and P-IR implementation of the efficient allocation and maximizes the market makers 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 are eventually cleared from the market without trading.[27] The optimal posted price mechanism does not coincide with the profit-maximizing efficient mechanism because of profit that is lost on efficient trades executed with a buyer's price of $p_B = 1 - \Delta_0$ or a seller's price of $p_S = \Delta_0$.

The Coase Theorem

We conclude this section by noting that our results uncover a dynamic connection between two fundamental theorems in economics, namely between the Coase Theorem, which states that, absent transaction costs, efficient trade of resource poses no problem, and the Myerson-Satterthwaite Theorem, which states that private information poses an insurmountable obstacle for efficient trade. In this section we proved that efficient trade is possible if storing enough traders is dynamically efficient, which shows that these theorems they can be reconciled by accounting for dynamics. Put differently, we have proven that the Coase Theorem applies if there is an appropriately designed dynamic market mechanism.

3.5 Applications

Our model and analysis also shed light on other questions of pertinent interest to economists such as in-house production and indirect taxes (and fees). In this section, we briefly discuss these applications. We also lay out how the optimal mechanism can be implemented in a prior-free, asymptotically optimal way.

[27]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 market 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.

In-House Production We start by considering an extension of the baseline model in which a profit-maximizing designer has in addition the ability to produce in-house at unit cost $\check{c}$. Alternatively, and equivalently, one can think of the market maker as being vertically integrated. Consider first the complete information benchmark. In this case, full surplus extraction is possible and the designer receives payoffs of $\overline{v} - \underline{c}$ and $\overline{v} - \overline{c}$ when trade is induced for efficient pairs and suboptimal pairs, respectively. The designer will therefore optimally produce in-house at times if and only if $\check{c} < \overline{c}$. Assume now that types are the agents' private information. In this case, the designer will optimally produce in-house at times if and only if $\check{c} < \Gamma(\overline{c})$. Since $\overline{c} < \Gamma(\overline{c})$, the designer will produce more often in-house when agent types are private information, reflecting the notion that agency costs foster integration (see e.g. Williamson, 1985). By the same token, under private information the designer will choose to produce in-house even when this is socially wasteful in the sense that $\overline{c} < \check{c} < \Gamma(\overline{c})$.[28] Interestingly, because $\Gamma(\overline{c})$ increases in the share p of efficient traders, the incentives for vertical integration increase in this share. That is, the larger is p the larger is the social waste the designer is willing to accept while still producing in-house. While the probability that in-house production occurs decreases in p, the threshold value for the in-house production cost of the profit-maximizing, vertically integrated market maker increases in p.

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

[28] Whether profitable vertical integration by a profit-maximizing market maker is indeed socially wasteful when compared to the equilibrium outcomes without integration (rather than to first-best) appears to depend on the size of $\check{c}$.

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 and is less than perfect 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.[29] This is analogous to the standard assumption in oligopoly models of indirect taxation that firms' profits can be observed and taxed. The analysis applies equally to discriminatory, uniform and fixed frequency market clearing.

Under a specific tax σ per unit traded with $\sigma > 0$, the relative value of a suboptimal trade compared to an efficient trade decreases from Δ_1 to $(\Delta_1 - \sigma)/(1 - \sigma)$. If $\sigma > \Delta_1$, this will induce the market maker to become perfectly patient. If $\sigma < \Delta_1$, by the results of Section 3, this will induce the maker to 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. 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.[30]

Prior-free, Asymptotically Optimal Implementation As noted, efficiency is not a distribution-free notion. A criticism of Bayesian mechanism design, often associated with Wilson (1987), is that it depends on the fine details of the environment such as assumptions about distributions and on higher-order beliefs, which, in reality, designers and possibly the agents may be uncertain or agnostic about.[31] In dynamic settings like ours, this creates a

[29] 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.

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

[31] In private values environments, this "Wilson" critique has led to the postulate that, for practicality, mechanisms endow the agents with dominant strategies and be free of the details of the underlying envi-

tension between efficiency and robustness as defined by Bergemann and Morris (2012). We now briefly sketch how this tension can be resolved or, at least, reduced by estimating the distributions from agents' reports, assuming for the remainder of this subsection that $\alpha = 0$. Exactly because efficiency is not a distribution-free concept, estimation proves useful even when the designer puts zero weight on profit.[32]

The basic idea is simple. With every additional trader arriving and reporting his type, an additional data point is generated, which, as time goes on and data accumulate, can be used to estimate the underlying distributional parameter p ever more precisely. For this purpose, it does not matter whether the designer and the agents are frequentists or Bayesian. The key is to maintain incentive compatibility for all agents and all types because otherwise reports induce biased estimates of the true distribution.[33] Maintaining incentive compatibility is possible though a little trickier than it may appear at first glance.

Suppose the designer shares the history of reports with the newly arriving agents. Even so, the agents' estimates of distributions will typically not be the same because the agents are privately informed about their own types. This requires imposing some leeway in the incentive compatibility constraints, which is possible because of the discrete type space. The designer can update the threshold $\tau^*(.)$ as efficient, non-identical suboptimal and inefficient pairs are cleared from the market upon arrival. Only suboptimal pairs which arrive to an empty market and identical suboptimal pairs are stored. By employing a first-in first-out queueing protocol, the designer can ensure that stored agents are unaffected if the market clearing policy is updated on the basis of their reports.

ronment. In turn, these have led to the development of literatures on robust mechanism design and on mechanisms that estimate distributions; see, respectively, Bergemann and Morris (2012) and Loertscher and Marx (2017) and the references therein.

[32] Most of the literature on mechanisms with estimation has focused on estimation to approximate Bayesian optimality (see, for example, Baliga and Vohra (2003); Segal (2003)). The only exception we are aware of that uses estimation in mechanism design with two-sided private information when the designer's objective is efficiency without a deficit is Loertscher and Mezzetti (2016), who use estimation to gauge aggregate demand and supply. Gershkov and Moldovanu (2009) consider achieving the efficient allocation in a dynamic environment with one-sided private information where the designer gradually learns the distribution of agents' values. In a more recent paper, Gershkov et al. (2017) consider a related problem in which the designer knows the distribution of agent types but must estimate the arrival process.

[33] In light of the strand of literature that suggests that design problems be approached by violating (or getting rid of) incentive compatibility constraints (see, for example, the discussion in Milgrom, 2007), it may be worth emphasizing that in dynamic settings like the present one this approach would prove problematic because absent incentive compatibility distributions cannot be learned from bids or reports.

4 Approximately Optimal Mechanisms

Many real-world market places do not allow for the flexible, discriminatory market clearing mechanisms that we have analyzed above. Some market places, such as electronic foreign exchange spot markets, clear markets uniformly but at a variable frequency while other exchanges have uniform market clearing that occurs at a fixed frequency. We now extend our analysis to these more restrictive market clearing mechanisms and then compare the performance of all our dynamic market clearing mechanisms with instantaneous market clearing.

4.1 Uniform and Fixed Frequency Market Clearing

Recall that under discriminatory market clearing, the market maker determines which agents are cleared from the market in each period. To expand our previous analysis analysis, we now introduce two additional notions of market clearing:

Definition 5. *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.*

We now derive the class of Bayesian optimal mechanisms under uniform and fixed frequency market clearing.

Uniform Market Clearing Under uniform market clearing, the state space, transition probabilities and reward function of the associated Markov decision process are the same as that of the Markov decision process derived in Section 3.2 for discriminatory market clearing. Uniform market clearing only affects the set of actions available to the mechanism designer in a given state. Let $\mathcal{A}'_{\boldsymbol{x}}$ denote the set of actions available to the mechanism 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 mechanism 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 $\langle \mathcal{X}, \mathcal{A}', P, r, \delta \rangle$.

As we did with discriminatory 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 6. *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 6. 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 *state* 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 quasi-birth-death process (see, for example, Latouche and Ramaswami (1999)). The transition matrix $\boldsymbol{P}'$ and stationary distribution of the Markov chain are computed in Appendix A.9. 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 2. *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 2 proceeds in a similar manner to the proof of Theorem 1, using a dynamic programming characterization of the optimal threshold τ^*. Algorithm 2 in Appendix B uses the optimal stopping condition derived from the Bellman equation to compute τ^*. An analogous result to Proposition 3, specifically that the designer's objective function is increasing in δ, p and Δ_α and τ^* is increasing in δ and decreasing in Δ_α, immediately follows from the dynamic programming characterization.[34]

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

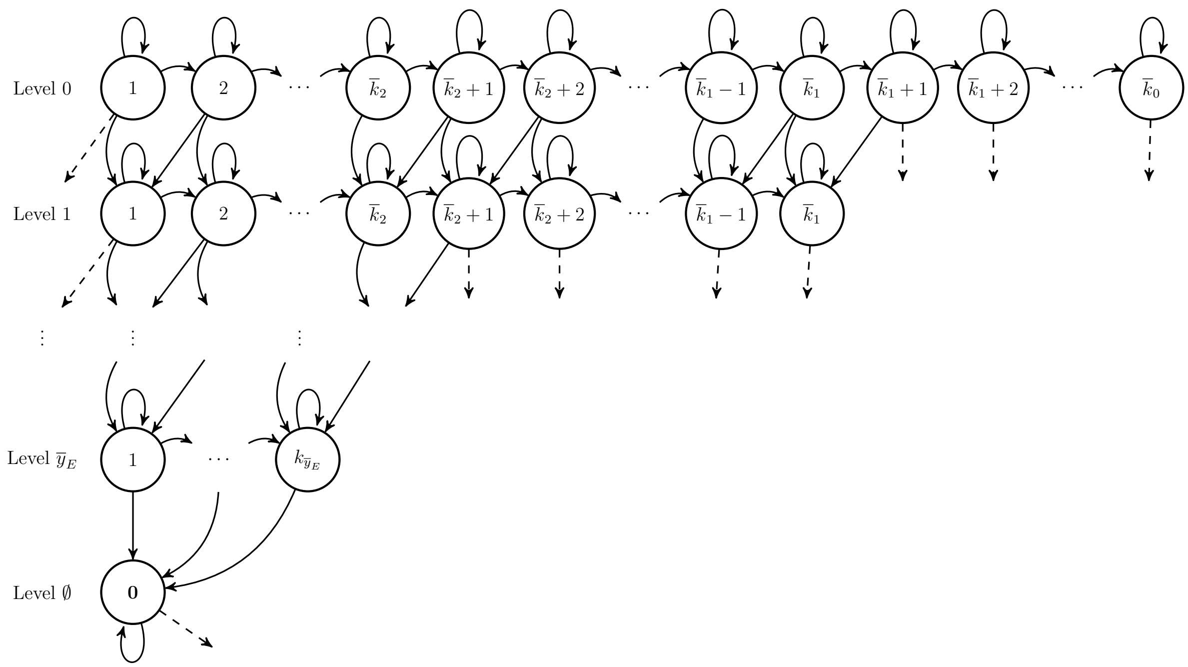

Figure 6: The structure of the quasi-birth-death under the threshold policy with threshold τ. Dashed arrows are used to denote some transitions to and from the state **0**.

Fixed Frequency Market Clearing Under fixed frequency market clearing, the state space of the limit 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 every τ periods, the market maker's expected discounted payoff is given by

$$W^F(\tau) = \frac{\delta^{\tau-1}}{1-\delta^\tau} \sum_{j=0}^{\tau} \sum_{k=0}^{\tau} \binom{\tau}{j} \binom{\tau}{k} (\min\{j,k\} + |j-k|\,\Delta_\alpha)\, p^{j+k}(1-p)^{2\tau-j-k}. \tag{16}$$

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 3, which can be found in Appendix B, uses this formula to compute the optimal market clearing threshold τ^*. Once again, we have that the designer's objective function is increasing in δ, p and Δ_α and τ^* is increasing in δ and decreasing in Δ_α. We also have the following result.

Corollary 4. *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.*

4.2 Dynamic vs. Instantaneous Market Clearing

Creating larger markets and employing increasingly sophisticated mechanisms may involve costs such as advertising and promotion, physical infrastructure investments, and labor. To assess the benefits from more sophisticated mechanisms, we now compare (in decreasing order of sophistication) discriminatory, uniform, fixed frequency market clearing and a period ex post efficient market mechanism that never stores trades and executes compatible trades instantaneously. Of course, as the designer uses an increasingly sophisticated dynamic mechanism, her payoff must increase since she solves a less constrained optimization problem. Furthermore, as δ increases, any of the dynamic mechanisms will significantly outperform instantaneous market clearing. However, we will also show that as δ increases, most gains from using a dynamic mechanism are reaped by clearing markets at a fixed (optimally chosen) frequency and there is little benefit associated with utilizing a state-dependent dynamic mechanism. Furthermore, we will show that, for sufficiently large discount factors, the social welfare gains under a given profit-maximizing dynamic market mechanism exceeds social welfare gains from using any less sophisticated ex post efficient mechanism. In other words, with sufficiently patient agents, a profit-maximizing but farsighted monopoly creates larger welfare gains than a benevolent but less sophisticated exchange.

Total expected discounted welfare gains period ex post efficient market mechanism, which for brevity we simply call *instantaneous trade*, are

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

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.[35]

Denote by $W^{D,\alpha}$ expected discounted welfare (starting from an empty market at $t = 0$) under discriminatory market clearing with a designer who maximizes $(1-\alpha)W + \alpha R$

[35] 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.

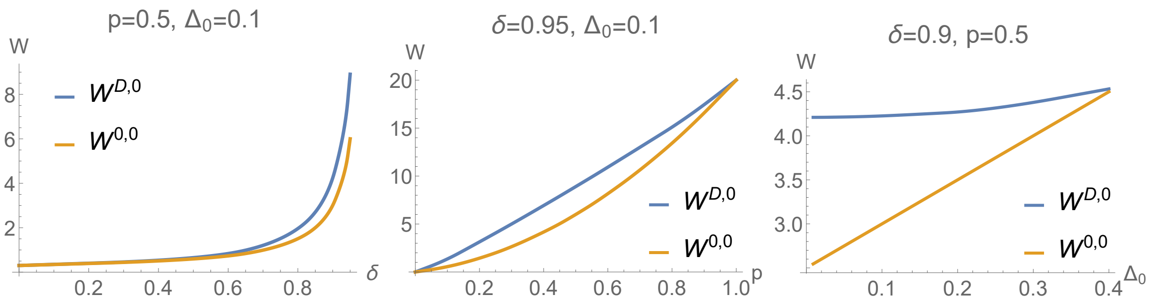

(a) δ comparative statics (b) p comparative statics (c) Δ_0 comparative statics

Figure 7: A numerical illustration of the performance of the optimal mechanism versus instantaneous trade for a variety of parameter values.

with $\alpha \in [0,1]$. Similarly, we use to notation $W^{U,\alpha}$, $W^{F,\alpha}$ and $W^{0,\alpha}$ for uniform market clearing, fixed frequency market clearing and instantaneous trade, respectively. Welfare gains under the optimal welfare-maximizing mechanism weakly exceed welfare gains under instantaneous trade as illustrated in Figure 7. 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 we consider impose less restrictions than fixed frequency market clear, 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 first-best outcome when $\delta = 0$.

Far less obvious is the comparison of social welfare gains under the optimal *profit-maximizing mechanism* to those under instantaneous trade. The following theorem contains this comparison, which is illustrated in Figure 8. Note that a simple argument based on less constrained optimization cannot be used to do this comparison. As one goes from periodically ex post efficient trade to profit-maximizing discriminatory market clearing, one not only eliminates constraints but also alters the objective that is maximized.

Theorem 3. *There exist $\underline{\delta}, \overline{\delta} \in [0,1)$ with $\underline{\delta} < \overline{\delta}$ such that*

- $W^{D,1}(\delta) \leq W^{0,0}(\delta)$ *for all* $\delta \leq \underline{\delta}$*,*
- $W^{D,1}(\delta) < W^{0,0}(\delta)$ *for all* $\delta \in (\underline{\delta}, \overline{\delta})$*,*
- $W^{D,1}(\overline{\delta}) = W^{0,0}(\overline{\delta})$ *and* $W^{D,1}(\delta) > W^{0,0}(\delta)$ *for all* $\delta > \overline{\delta}$*.*

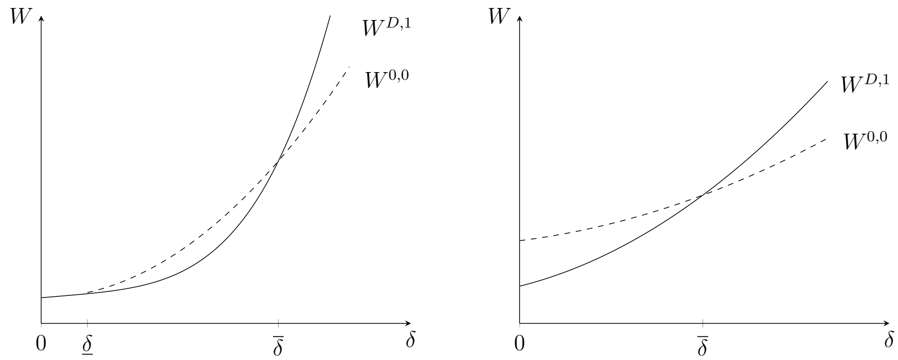

(a) $\Delta_1 > 0$ (b) $\Delta_1 \leq 0$

Figure 8: Illustration of Theorem 3.

Theorem 3 effectively says that for δ sufficiently large, welfare under profit-maximizing discriminatory market clearing exceeds welfare under myopically, periodically ex post efficient instantaneous trade. A natural, and in our view not too far-fetched, interpretation is that profit-maximizing marker maker is an Internet giant while the period 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 prove 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.

To develop the intuition behind this result, we start with the case $\Delta_1 > 0$. Here, no traders are stored under any mechanism for small values of δ. Thus, welfare under instantaneous trade is equal to welfare under the profit-maximizing mechanism. For intermediate values of δ, the number of traders stored under the dynamically efficient mechanism is closer to zero than to the larger number of trades stored under the profit-maximization mechanism. Consequently, $W^{D,1}(\delta) < W^{0,0}(\delta)$ for such values of δ. As $\delta \to 1$ both the welfare-maximizing and profit-maximizing designer become perfectly patient and so for sufficiently large δ, welfare under the profit-maximizing mechanism exceeds welfare under myopically, periodically ex post efficient instantaneous trade. The intuition for the $\Delta_1 \leq 0$ case is analogous, except that an inefficiently large number of traders are stored under profit-maximizing discriminatory market clearing even for $\delta = 0$.

We now compare the benefits from increasing sophistication for a given α as δ goes to 1.

Theorem 4 (Gains from Sophistication)**.** *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. Formally, we have*

$$\lim_{\delta \to 1} \frac{W^{D,\alpha}(\delta) - W^{U,\alpha}(\delta)}{W^{D,\alpha}(\delta)} = 0 = \lim_{\delta \to 1} \frac{W^{U,\alpha}(\delta) - W^{F,\alpha}(\delta)}{W^{U,\alpha}(\delta)} = 0$$

and

$$\lim_{\delta \to 1} \frac{W^{F,\alpha}(\delta) - W^{0,\alpha}(\delta)}{W^{F,\alpha}(\delta)} = (1-p)(1-2\Delta_\alpha).$$

The absolute gains from a higher degree of sophistication are increasing in δ *and diverge as* δ *approaches 1. Formally, the functions* $W^{D,\alpha}(\delta) - W^{U,\alpha}(\delta)$*,* $W^{U,\alpha}(\delta) - W^{F,\alpha}(\delta)$ *and* $W^{F,\alpha}(\delta) - W^{0,\alpha}(\delta)$ *with domain* $(0,1)$ *are positive, increasing in* δ *and diverge as* $\delta \to 1$*.*

Theorem 4 describes absolute and relative gains from (additional) sophistication and is illustrated in Figures 9 and 10. As δ becomes sufficiently large, the absolute gains from additional sophistication – e.g. from going from fixed frequency to uniform market clearing or from uniform to discriminatory market clearing – increase. Although the functions $W^{D,\alpha}(\delta)$, $W^{U,\alpha}(\delta)$, $W^{F,\alpha}(\delta)$ converge to the same limit, the absolute gains diverge on $(0,1)$ as $\delta \to 1$ because of the differences in the speed of convergence. If additional sophistication came at a fixed cost that is unrelated to δ, the first part of Theorem 4 would provide a powerful argument in favor of additional sophistication and sophistication. However, the cost of additional sophistication may plausibly be proportional to δ because additional sophistication comes at the cost of requiring additional investments (in infrastructure) that scales with the number of traders stored (and thus, indirectly, with δ). In this case, the relative gains are the relevant benchmark and the gist of Theorem 4 is that the key is to depart from instantaneous trade in favor of even the most moderate form of sophistication (i.e. fixed frequency market clearing).

Another pertinent issue in the design of two-sided markets is the need to "bring both sides of the market on board" (see e.g. 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 corollary of Theorem 3 sheds new light on this question.

Corollary 5. *There exist* $\delta_D, \delta_U, \delta_F \in (0,1)$ *such that* $W^{D,1}(\delta) > W^{U,0}(\delta)$ *for all* $\delta > \delta_D$*,* $W^{U,1}(\delta) > W^{F,0}(\delta)$ *for all* $\delta > \delta_U$ *and* $W^{F,1}(\delta) > W^{0,0}(\delta)$ *for all* $\delta > \delta_F$*.*

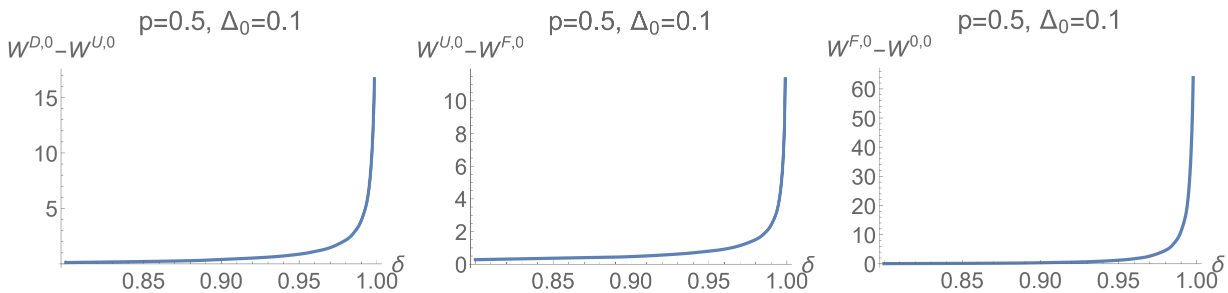

(a) Discrim. vs. uniform (b) Uniform vs. fixed freq. (c) Fixed freq. vs. decent.

Figure 9: The absolute gains of higher degrees of sophistication.

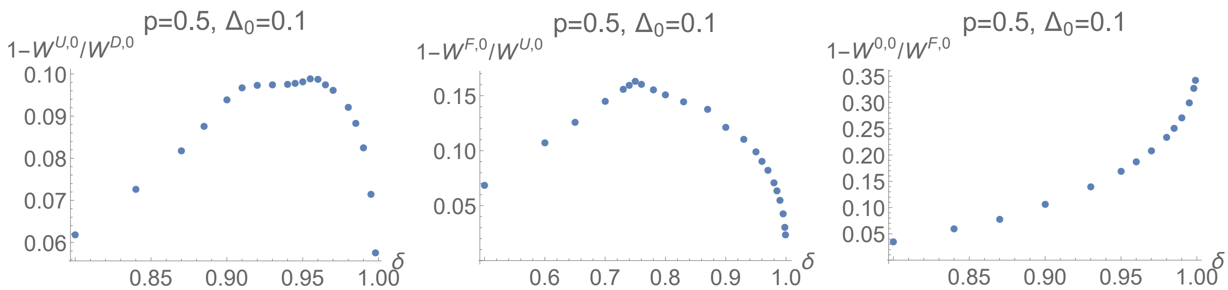

(a) Discrim. vs. uniform (b) Uniform vs. fixed freq. (c) Fixed freq. vs. decent.

Figure 10: The relative gains of additional sophistication.

In words, Corollary 5 says that for sufficiently large discount factors, a profit-maximizing platform will generate more welfare than an otherwise welfare-maximizing platform using a less sophisticated market clearing mechanism. 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, say, a welfare-maximizing platform with a lower degree of sophistication, by Corollary 5 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.

5 Extensions

We now discuss how our methodology and main results generalize in a variety of ways, with further extensions included in Appendix C. Although a threshold policy will not generally be optimal, in each of the extensions we consider in this paper threshold policies can be used to partition the state space so that computing the optimal policy becomes tractable. We develop this methodology in detail in Appendix D.

5.1 General Arrival Processes

Under discriminatory market clearing, the assumption that buyers and sellers do not arrive in pairs can easily be relaxed. To see this, suppose 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.[36] 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 3.2. Under fixed frequency market clearing, the optimal policy can be computed similarly by modifying (16).

With uniform market clearing (see Definition 5), dealing with unpaired agents is more problematic because from the market maker's perspective, no immediate reward is earned when unpaired agents are cleared from the market but they are useful for forming pairs in the future. Thus, there is no upper bound on the number of unpaired agents stored under the optimal policy and allowing for unpaired agents significantly increases the size of the state space of the associated Markov decision process. We deal with this case in Appendix D.

Our results immediately generalize to the case in which pairs of buyers and sellers arrive according to a Poisson process.[37] Under every extension of the arrival process considered in this paper, Corollary 1 and the vertical integration and taxation results of Section 3.5 generalize because they depend directly on Δ_α but not on the arrival process. Furthermore, once the state space and transition probabilities are determined for a given arrival process, the results regarding the degree of market sophistication in Corollary 5 also generalize. These results depend directly on the constraints applied to the optimization problem and by virtue of the properties of the Poisson equation, the proof of Lemma 3 in Appendix A.13 applies

[36]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.

[37]In Appendix C we further generalize this to allow for more general renewal processes. In Appendix D we allow for buyers and sellers to arrive according to independent Poisson processes.

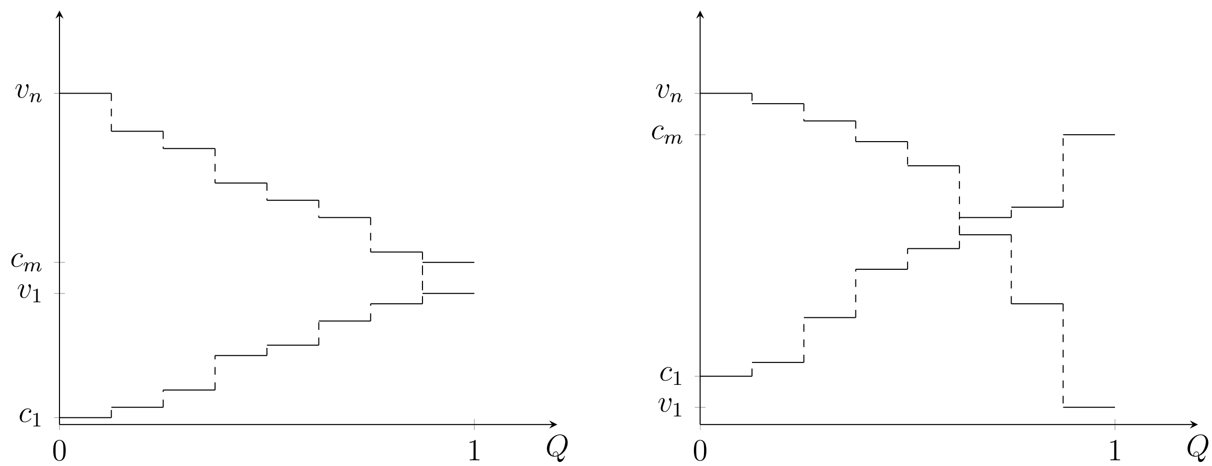

(a) $v_2 > c_m > v_1 > c_{m-1}$ holds

(b) $v_2 > c_m > v_1 > c_{m-1}$ violated

Figure 11: In order to consider general discrete type spaces, we must deal with two cases.

to any arrival process.

5.2 General Type Spaces

Our analysis also extends to models with richer, discrete type spaces. Assume that buyers draw their types independently from a discrete distribution F whose support is given by $\{v_1, \ldots, v_n\}$ with $v_1 < \cdots < v_n$ for $n \in \mathbb{N}$ and that sellers draw their types independently from a discrete distribution G with support $\{c_1, \ldots, c_m\}$, where $c_1 < \cdots < c_m$ for $m \in \mathbb{N}$. For $i \in \{1, \ldots, n-1\}$ and $j \in \{2, \ldots, m\}$, the virtual type functions are

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

while for $i = n$ and $j = 1$, they are $\Phi(v_n) = v_n$ and $\Gamma(c_1) = c_1$.

Our methodology generalizes immediately under the assumption $v_2 > c_m > v_1 > c_{m-1}$, which serves the same purpose as the restriction $\overline{v} > \overline{c} > \underline{v} > \underline{c}$ in the binary type setup (see Panel (a), Figure 11). One just has to expand the state space of the Markov decision process so that it includes each type of efficient and suboptimal trade. The main challenge in dealing with general type spaces such that $v_2 > c_m > v_1 > c_{m-1}$ does not hold is that the market maker will now optimally store some types of infeasible trades which could prove useful for rematching in future periods (see Panel (b), Figure 11). Since there is no cost associated with storing such trades there is no bound on the number that can accumulate under the optimal policy. Thus, the existence of this type of infeasible trade substantially increases the size of the state space of the underlying Markov decision process. However, this can be dealt with by applying the same methodology used to solve the problem created by unpaired agents.[38]

For our main comparative statics results (specifically Corollary 1 and Corollary 5, which compare outcomes under market makers with differing objectives) to generalize, a new condition, which we call *dynamic regularity*, is sufficient whereas Myerson's regularity condition is not. Distributions F and G are said to satisfy *dynamic regularity* if Φ and Γ are non-decreasing and if, for $i \in \{1, \ldots, n-1\}$ and $j \in \{2, \ldots, m\}$,

$$\Phi(v_{i+1}) - \Phi(v_i) > v_{i+1} - v_i \quad \text{and} \quad \Gamma(c_j) - \Gamma(c_{j-1}) > c_j - c_{j-1} \tag{18}$$

holds. 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. Like in static settings, Myerson's regularity condition is sufficient for pointwise maximization to be incentive compatible. 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. Under this regularity assumption, we also obtain the following generalization of the possibility result from Section 3.4.

Proposition 6. *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, for $\alpha = 0$, the expected discounted profit of the designer under the optimal mechanism is positive.*

[38] That is, analogous to the methodology developed in Appendix D, we consider a related Markov decision process in which any infeasible trades containing an efficient agent are replaced by a trade containing the efficient agent paired with the least efficient agent that creates a feasible trade. By determining the optimal policy of the modified Markov decision process, we can identify a finite number of candidate optimal policies for the original Markov decision process.

Lastly, if we maintain the assumption of binary types, our analysis, including the mechanism design approach of Section 3.5, also generalizes to agents with multi-unit demands and multi-unit capacities as discussed in Appendix C.3. The key insight is that the design problem retains the property that agents' private information pertains to a single dimension, namely the number of units a buyer values highly respectively the number of units a seller can produce at low costs.

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 microfoundations 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 from waiting to clear the market against the opportunity cost of delay.

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 (i.e. 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 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 here: 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 trade 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 social welfare gains under all of these mechanisms converge and diverge from the gains under the periodic ex post efficient mechanism that never stores trade.

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 ride-sharing service provides 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 has 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.[39] In light of this, an interesting extensionof 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 interdependnet 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 given environment.

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

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*, forthcoming.

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.

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

Bergemann, D. and S. Morris (2012): *Robust Mechanism Design*, Hackensack, NJ: World Scientific Press.

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

Blackwell, D. (1965): "Discounted Dynamic Programming," *Annals of Mathematics and Statistics*, 36, 226–235.

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.

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.

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 (2009): “Learning about the Future and Dynamic Efficiency,” *The American Economic Review*, 99, 1576–1587.

——— (2010): “Efficient Sequential Assignment with Incomplete Information,” *Games and Economic Behavior*, 68, 144–154.

Gershkov, A., B. Moldovanu, and P. Strack (2017): “Revenue Maximizing Mechanisms with Strategic Customers and Unknown, Markovian Demand,” *Management Science*, forthcoming.

Gittins, J. C. (1979): “Bandit Processes and Dynamic Allocation Indices,” *Journal of the*

Royal Statistical Society, Series B, 41, 148–177.

Glynn, P. W. and S. P. Meyn (1996): "A Liapounov Bound for Solutions of the Poisson Equation," *The Annals of Probability*, 24, 916–931.

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.

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

Howard, R. A. (1960): *Dynamic Programming and Markov Decision Processes*, Cambridge, Massachusetts: The M.I.T Press.

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.

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.

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., L. Marx, and T. Wilkening (2015): "A Long Way Coming: Designing Centralized Markets with Privately Informed Buyers and Sellers," *Journal of Economic Literature*, 53, 857–897.

Loertscher, S. and C. Mezzetti (2016): "Dominant Strategy, Double Clock Auctions with Estimation-Based Tâtonnement," Working paper.

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

Malamud, S. and M. Rostek (2016): "Decentralized Exchange," 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.

Milgrom, P. R. (2007): "Package Auctions and Exchanges," *Econometrica*, 75, 935–965.

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.

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.

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

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.

Stokey, N. L., R. E. Lucas, and E. C. Prescott (1989): *Recursive Methods in Economic Dynamics*, Cambridge: Harvard University Press.

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.

Williamson, O. (1985): *The Economic Institutions of Capitalism*, New York: The Free Press.

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

Yavaş, A. (1996): "Matching of Buyers and Sellers by Brokers: A Comparison of Alternative Commission Structures," *Real Estate Economics*, 24, 97–112.

Appendices

A Proofs

A.1 Proof of Proposition 1

Proof. We start by proving the result with Bayesian incentive compatibility and individual rationality constraints. When B_t reports $\hat{v}_t$ and S_t reports $\hat{c}_t$ the respective interim discounted allocation probabilities, assuming all other agents arriving after period $t-1$ report truthfully, are given by

$$\begin{aligned} q^{B_t}(\hat{v}_t) &= \sum_{i=t}^{\infty} \sum_{\boldsymbol{h}_i \in \mathcal{H}_i} \delta^{i-1} Q_i^{B_t}(\boldsymbol{h}_i) \mathbb{P}(\boldsymbol{H}_i = \boldsymbol{h}_i | V_t = \hat{v}_t), \\ q^{S_t}(\hat{c}_t) &= \sum_{i=t}^{\infty} \sum_{\boldsymbol{h}_i \in \mathcal{H}_i} \delta^{i-1} Q_i^{S_t}(\boldsymbol{h}_i) \mathbb{P}(\boldsymbol{H}_i = \boldsymbol{h}_i | C_t = \hat{c}_t), \end{aligned} \tag{19}$$

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

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

Denote the respective probability mass functions of V_t and C_t by f and g. Using the fact that the individual rationality constraints bind for the worst-off types and (4) we have

$$\begin{aligned} R &= \sum_{i=1}^{\infty} \boldsymbol{E}\left[m^{B_i}(V_i) - m^{S_i}(C_i)\right] \\ &= \sum_{i=1}^{\infty} \Big\{ \left[\underline{v} q^{B_i}(\underline{v}) - \underline{c}(q^{S_i}(\underline{c}) - q^{S_i}(\overline{c})) - \overline{c} q^{S_i}(\overline{c})\right] f(\underline{v}) g(\underline{c}) \\ &+ \left[\overline{v}(q^{B_i}(\overline{v}) - q^{B_i}(\underline{v})) + \underline{v} q^{B_i}(\underline{v}) - \underline{c}(q^{S_i}(\underline{c}) - q^{S_i}(\overline{c})) - \overline{c} q^{S_i}(\overline{c})\right] f(\overline{v}) g(\underline{c}) \\ &+ \left[\underline{v} q^{B_i}(\underline{v}) - \overline{c} q^{S_i}(\overline{c})\right] f(\underline{v}) g(\overline{c}) \\ &+ \left[\overline{v}(q^{B_i}(\overline{v}) - q^{B_i}(\underline{v})) + \underline{v} q^{B_i}(\underline{v}) - \overline{c} q^{S_i}(\overline{c})\right] f(\overline{v}) g(\overline{c}) \Big\}. \end{aligned}$$

Rearranging, using that $f(\overline{v}) = g(\underline{c}) = p$ and $f(\underline{v}) = g(\overline{c}) = 1 - p$ and making the substitutions from (6) we obtain

$$\begin{aligned}
R = \sum_{i=1}^{\infty} \Bigg\{ & \left[q^{B_i}(\underline{v}) \left(\underline{v} + \frac{f(\overline{v})(\underline{v} - \overline{v})}{f(\underline{v})} \right) - \underline{c} q^{S_i}(\underline{c}) \right] f(\underline{v}) g(\underline{c}) \\
& + \left[\overline{v} q^{B_i}(\overline{v}) - \underline{c} q^{S_i}(\underline{c}) \right] f(\overline{v}) g(\underline{c}) \\
& + \left[q^{B_i}(\underline{v}) \left(\underline{v} + \frac{f(\overline{v})(\underline{v} - \overline{v})}{f(\underline{v})} \right) - q^{S_i}(\overline{c}) \left(\overline{c} + \frac{g(\underline{c})(\overline{c} - \underline{c})}{g(\overline{c})} \right) \right] f(\underline{v}) g(\overline{c}) \\
& + \left[\overline{v} q^{B_i}(\overline{v}) - q^{S_i}(\overline{c}) \left(\overline{c} + \frac{g(\underline{c})(\overline{c} - \underline{c})}{g(\overline{c})} \right) \right] f(\overline{v}) g(\overline{c}) \Bigg\} \\
= & \sum_{i=1}^{\infty} \sum_{v_i=\underline{v}}^{\overline{v}} \sum_{c_i=\underline{c}}^{\overline{c}} \left[\Phi(v_i) q^{B_i}(v_i) - \Gamma(c_i) q^{S_i}(c_i) \right] f(v_i) g(c_i) \\
= & \sum_{i=1}^{\infty} \sum_{v_i=\underline{v}}^{\overline{v}} \Phi(v_i) q^{B_i}(v_i) f(v_i) - \sum_{i=1}^{\infty} \sum_{c_i=\underline{c}}^{\overline{c}} \Gamma(c_i) q^{S_i}(c_i) g(c_i).
\end{aligned}$$

Finally, using (19) gives

$$\begin{aligned}
R = & \sum_{i=1}^{\infty} \sum_{t=i}^{\infty} \sum_{h_t \in \mathcal{H}_t} \sum_{v_i=\underline{v}}^{\overline{v}} \delta^{t-1} \Phi(v_i) Q_t^{B_i}(h_t) \mathbb{P}(H_t = h_t | V_i = v_i) f(v_i) \\
& - \sum_{i=1}^{\infty} \sum_{t=i}^{\infty} \sum_{h_t \in \mathcal{H}_t} \sum_{c_i=\underline{c}}^{\overline{c}} \delta^{t-1} \Gamma(c_i) Q_t^{S_i}(h_t) \mathbb{P}(H_t = h_t | C_i = c_i) g(c_i) \\
= & \sum_{i=1}^{\infty} \sum_{t=i}^{\infty} \sum_{h_t \in \mathcal{H}_t} \delta^{t-1} \Phi(v_i) Q_t^{B_i}(h_t) \mathbb{P}(H_t = h_t) - \sum_{i=1}^{\infty} \sum_{t=i}^{\infty} \sum_{h_t \in \mathcal{H}_t} \delta^{t-1} \Gamma(c_i) Q_t^{S_i}(h_t) \mathbb{P}(H_t = h_t) \\
= & \sum_{t=1}^{\infty} \sum_{i=1}^{t} \sum_{h_t \in \mathcal{H}_t} \delta^{t-1} \left(\Phi(v_i) Q_t^{B_i}(h_t) - \Gamma(c_i) Q_t^{S_i}(h_t) \right) \mathbb{P}(H_t = h_t)
\end{aligned}$$

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

$$R = \sum_{i=1}^{\infty} \boldsymbol{E}_{V_i, C_i, \boldsymbol{H}_{t-1}} \left[m^{B_i}(V_i, \boldsymbol{H}_{t-1}) - m^{S_i}(C_i, \boldsymbol{H}_{t-1}) \right].$$

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

$$R = \sum_{i=1}^{\infty} \boldsymbol{E}_{V_i, C_i, \boldsymbol{H}_{t-1}} \left[m^{B_i}(V_i, C_i, \boldsymbol{H}_{t-1}) - m^{S_i}(C_i, V_i, \boldsymbol{H}_{t-1}) \right].$$

□

A.2 Proof of Theorem 1

Proof. The proof proceeds by deriving the optimal policy of the Markov decision policy and mapping this to the optimal allocation rule, before checking the relevant incentive constraints. The optimal policy must immediately clear all efficient pairs. Since the arrival process is stationary and the discount factor is constant, the optimal policy is stationary. 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 by stationarity, if x_S suboptimal trades are stored under the 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. 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 benefit of storing an additional suboptimal pair converges to zero, while the benefit of immediately clearing a suboptimal pair is Δ_α. Therefore, there exists a maximum number τ^* of suboptimal trades which can be optimally stored. Thus, 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 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.[40] 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

$$q^{B_t}(\overline{v}) = \delta^{t-1} \left[p + (1-p) q(\overline{v}, \overline{c}) \right]. \tag{20}$$

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

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

$$q^{B_t}(\underline{v}) = \delta^{t-1} p q(\underline{v}, \underline{c}). \tag{21}$$

Comparing (20) and (21) 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 Proposition 2

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

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

where $\lambda = p(1-p)$. The stationary distribution $\boldsymbol{\kappa}$ satisfies $\boldsymbol{\kappa P} = \boldsymbol{\kappa}$ and we solve

$$\sum_{i=0}^{\tau^*} \kappa_i = 1, \quad \kappa_0 = \kappa_0(1-2\lambda) + \kappa_1\lambda, \quad \kappa_1 = 2\kappa_0\lambda + \kappa_1(1-2\lambda) + \kappa_2\lambda,$$
$$\forall i \in \{1, \dots, \tau^*-1\}, \quad \kappa_i = \kappa_{i-1}\lambda + \kappa_i(1-2\lambda) + \kappa_{i+1}\lambda, \quad \kappa_{\tau^*} = \kappa_{\tau^*-1}\lambda + \kappa_{\tau^*}(1-\lambda).$$

The last equation is a second order difference equation with constant coefficients. We substitute a trial solution $\kappa_i = \psi^i > 0$ into this equation and solve the characteristic equation of the recursion. This gives

$$\psi^i = \psi^{i-1}\lambda + \psi^i(1-2\lambda) + \psi^{i+1}\lambda \quad \Rightarrow \psi^2 - 2\psi + 1 = 0 \quad \Rightarrow \psi = 1,$$

so the general solution to the difference equation is $\kappa_i = c_1 + c_2 i$. Here $c_1, c_2 \in \mathbb{R}$ are unknown constants. Substituting the general solution into the equations for κ_{τ^*} we obtain $\kappa_{\tau^*} = \kappa_{\tau^*-1}$. Substituting this into the equation for κ_{τ^*-1} gives $\kappa_{\tau^*-1} = \kappa_{\tau^*-2}$. Thus we require $c_1 + (\tau^*-1)c_2 = c_1 + (\tau^*-2)c_2$ which implies that $c_2 = 0$. Substituting $\kappa_i = c_1$ into the equations for κ_0 and κ_1 and solving these equations simultaneously gives $\kappa_0 = c_1/2$ and $\kappa_1 = c_1$. Finally, using the normalization equation, we have

$$\frac{c_1}{2} + \sum_{i=1}^{\tau^*} c_1 = 1 \quad \Rightarrow c_1 + 2c_1\tau^* = 2 \quad \Rightarrow c_1 = \frac{2}{2\tau^*+1}.$$

Thus, 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 and $(1-p)^2$ a $(\overline{v},\underline{c})$ pair and a $(\underline{v},\overline{c})$ pair arrive respectively, 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

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

The expression for $W_t^{D,\alpha}(\tau)$ 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, which 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. □

A.4 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 market 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

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

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

Next, examining (12), it can be seen that for $x_S \in \{0, 1, \ldots, \tau^*\}$ an increase in Δ_α and δ 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 is increasing in Δ_α and δ. For $x_S \in \{1, \ldots, \tau^*-1\}$, ranking the outcomes on the right-hand side of (12) 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 (i.e. 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. □

A.5 Proof of Corollary 1

Proof. Differentiating Δ_α yields

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

Recall 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.6 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 $\bar{c}$ receive a payoff of zero whenever they report truthfully, since under truthful reporting these agents will only accept prices of $\underline{v}$ and $\bar{c}$ 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 be cleared from 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 are no balanced budget posted prices $p_S = p_S$ that can clear efficient trades and both types of suboptimal trades. □

A.7 Proof of Corollary 2

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.8 Proof of Proposition 5

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 from 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 buying 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.9 The Markov Transition Matrix for Uniform Market Clearing

The transition matrix $\boldsymbol{P}'$ of the order book Markov chain $\{\boldsymbol{Y}_t\}_{t\in\mathbb{N}}$ under uniform market clearing has a block structure of the form

$$\boldsymbol{P}' = \begin{pmatrix} \boldsymbol{A}_{\emptyset\emptyset} & \boldsymbol{A}_{\emptyset 0} & \boldsymbol{0} & \boldsymbol{0} & \cdots & \boldsymbol{0} \\ \boldsymbol{A}_{0\emptyset} & \boldsymbol{A}_{00} & \boldsymbol{A}_{01} & \boldsymbol{0} & \cdots & \boldsymbol{0} \\ \boldsymbol{A}_{1\emptyset} & \boldsymbol{0} & \boldsymbol{A}_{11} & \boldsymbol{A}_{12} & \ddots & \boldsymbol{0} \\ \boldsymbol{A}_{2\emptyset} & \boldsymbol{0} & \boldsymbol{0} & \boldsymbol{A}_{22} & \ddots & \boldsymbol{0} \\ \vdots & \vdots & \ddots & \ddots & \ddots & \ddots \\ \boldsymbol{A}_{\overline{x}_E\emptyset} & \boldsymbol{0} & \boldsymbol{0} & \boldsymbol{0} & \ddots & \boldsymbol{A}_{\overline{x}_E\overline{x}_E} \end{pmatrix}.$$

Take $i \in \{0, 1, \ldots, \overline{x}_E\}$. The $\boldsymbol{A}_{i\emptyset}$ blocks are $\overline{k}_i \times 1$ matrices which encode the transitions from level i to the market clearing state. The $\boldsymbol{A}_{ii}$ blocks are $\overline{k}_i \times \overline{k}_i$ matrices which encode the transitions from level i to level i. The $\boldsymbol{A}_{ii+1}$ blocks are $\overline{k}_i \times \overline{k}_{i+1}$ matrices which encode the transitions from level i to level $i+1$. Letting $\mu_0 = (1-p)^2$ and $\mu_1 = p(1-p)$ we have $\boldsymbol{A}_{i\emptyset} = (\mu_1, \boldsymbol{0}, \boldsymbol{A}_i)'$, where $\boldsymbol{A}_i = (p^2, p^2+\mu_1, \ldots, p^2+\mu_1, p^2+2\mu_1)'$ is a vector of length $\overline{k}_i - \overline{k}_{i+1}$. Furthermore,

$$\boldsymbol{A}_{ii} = \begin{pmatrix} \mu_0 & \mu_1 & \ddots & 0 \\ \ddots & \ddots & \ddots & \ddots \\ 0 & 0 & \ddots & \mu_1 \\ 0 & 0 & \ddots & \mu_0 \end{pmatrix} \quad \text{and} \quad \boldsymbol{A}_{ii+1} = \begin{pmatrix} p^2 & 0 & \cdots & 0 \\ \mu_1 & p^2 & \ddots & 0 \\ \ddots & \ddots & \ddots & \ddots \\ 0 & 0 & \ddots & p^2 \\ 0 & 0 & \ddots & \mu_1 \\ \vdots & \ddots & \ddots & \ddots \\ 0 & 0 & \ddots & 0 \end{pmatrix}.$$

Finally $\boldsymbol{A}_{\overline{x}_E,\emptyset} = (p^2+\mu_1, \ldots, p^2+\mu_1, p^2+2\mu_1)'$ is a $\overline{k}_{\overline{x}_E} \times 1$ matrix, $\boldsymbol{A}_{\emptyset 0} = \begin{pmatrix} 2\mu_1 & \boldsymbol{0} \end{pmatrix}$ is a $1 \times \tau$ matrix and $\boldsymbol{A}_{\emptyset\emptyset} = \begin{pmatrix} p^2 + \mu_0 \end{pmatrix}$.

From Figure 6 and the structure of $\boldsymbol{P}'$, it can be seen that the Markov chain $\{Y_t\}_{t\in\mathbb{N}}$ is similar in nature to a level-dependent quasi-birth-and-death process (see, for example, Latouche and Ramaswami (1999)). The stationary distribution $\boldsymbol{\pi}$ satisfies $\boldsymbol{\pi}\boldsymbol{P}' = \boldsymbol{\pi}$. That is, we must solve the system of equations

$$\boldsymbol{\pi}_\emptyset + \sum_{i=0}^{\overline{x}_E} \boldsymbol{\pi}_i \cdot \boldsymbol{e} = 1, \quad \boldsymbol{\pi}_0 \boldsymbol{A}_{00} + \boldsymbol{\pi}_\emptyset \boldsymbol{A}_{\emptyset 0} = \boldsymbol{\pi}_0,$$
$$\forall i \in \{1, \ldots, \overline{x}_E\}, \quad \boldsymbol{\pi}_{i-1} \boldsymbol{A}_{i-1i} + \boldsymbol{\pi}_i \boldsymbol{A}_{ii} = \boldsymbol{\pi}_i.$$

Here, $\boldsymbol{e}$ is a vector of ones of the appropriate length. Since, the matrix $\boldsymbol{I} - \boldsymbol{A}_{ii}$ is invertible we have

$$\boldsymbol{\pi}_0 = \boldsymbol{\pi}_\emptyset \boldsymbol{A}_{\emptyset 0}(\boldsymbol{I} - \boldsymbol{A}_{00})^{-1} \quad \text{and} \quad \boldsymbol{\pi}_i = \boldsymbol{\pi}_{i-1}\boldsymbol{A}_{i-1i}(\boldsymbol{I} - \boldsymbol{A}_{ii})^{-1}.$$

Thus, for all $i \in \{0, 1, \ldots, \overline{x}_E\}$, we have

$$\boldsymbol{\pi}_i = \boldsymbol{\pi}_\emptyset \boldsymbol{A}_{\emptyset 0}(\boldsymbol{I} - \boldsymbol{A}_{00})^{-1} \prod_{j=1}^{i} \boldsymbol{A}_{j-1j}(\boldsymbol{I} - \boldsymbol{A}_{jj})^{-1}, \quad \text{with} \quad \boldsymbol{\pi}_\emptyset = 1 - \sum_{i=0}^{\tau} \boldsymbol{\pi}_i \cdot \boldsymbol{e}.$$

Note that the matrix $\boldsymbol{A}_{i-1i}(\boldsymbol{I} - \boldsymbol{A}_{ii})^{-1}$ has a probabilistic interpretation. The (j,k)th entry of this matrix is the expected number of visits to state k in level $i-1$ before the process moves to level i, given that the matrix started in state j in level $i-1$.

Let $\boldsymbol{V}_\tau^U(i)$ denote the vector of the expected values associated with being in each state in level i under the threshold policy characterized by τ. Then using $\boldsymbol{P}'$, the linear system defined in (29) and (30) can be rewritten in matrix form as

$$(\boldsymbol{I} - \delta \boldsymbol{P}') \begin{pmatrix} \boldsymbol{V}_\tau^U(0) \\ \boldsymbol{V}_\tau^U(1) \\ \vdots \\ \boldsymbol{V}_\tau^U(\overline{x}_E) \\ \boldsymbol{V}_\tau^U(\boldsymbol{0}) \end{pmatrix} = \delta \begin{pmatrix} \boldsymbol{r}'_\tau(0) \\ \boldsymbol{r}'_\tau(1) \\ \vdots \\ \boldsymbol{r}'_\tau(\overline{x}_E) \\ \boldsymbol{r}'_\tau(\boldsymbol{0}) \end{pmatrix}.$$

Here, $\boldsymbol{r}'_\tau(\boldsymbol{0}) = p^2$ and for $i \in \{0, \ldots, \overline{x}_E - 1\}$ we have $\boldsymbol{r}'_\tau(i) = (\mu_1(1+i), \boldsymbol{0}, \boldsymbol{r}'_i)'$ where

$$\boldsymbol{r}'_i = \begin{pmatrix} p^2(i + 1 + \Delta_\alpha(\overline{k}_{i+1} + 1)) \\ p^2(i + 1 + \Delta_\alpha(\overline{k}_{i+1} + 2)) + \mu_1(i + 1 + \Delta_\alpha(\overline{k}_{i+1} + 1)) \\ \vdots \\ p^2(i + 1 + \Delta_\alpha(\overline{x}_I - i\overline{k} - 1)) + \mu_1(i + 1 + \Delta_\alpha(\overline{x}_I - \overline{k}_i - 2)) \\ p^2(i + 1 + \Delta_\alpha \overline{k}_i) + 2\mu_1(i + 1 + \Delta_\alpha(\overline{k}_i - 1)) + \mu_1(1 + 2\Delta_\alpha) \end{pmatrix},$$

is a $\overline{k} \times 1$ column vector and

$$\boldsymbol{r}'_\tau(\overline{x}_E) = \begin{pmatrix} p^2(1 + \overline{x}_E + \Delta_\alpha) + \mu_1(1 + \overline{x}_E) \\ p^2(1 + \overline{x}_E + 2\Delta_\alpha) + \mu_1(1 + \overline{x}_E + \Delta_\alpha) \\ \vdots \\ p^2(1 + \overline{x}_E + \Delta_\alpha(\overline{k}_{\overline{x}_E} - 1)) + \mu_1(1 + \overline{x}_E + \Delta_\alpha(\overline{k}_{\overline{x}_E} - 2)) \\ p^2(1 + \overline{x}_E + \overline{k}_{\overline{x}_E}) + 2\mu_1(\overline{x}_E + \Delta_\alpha(\overline{k}_{\overline{x}_E} - 1)) + \mu_1(1 + 2\Delta_\alpha) \end{pmatrix}.$$

A.10 Proof of Theorem 2

Proof. If we use the notation of Blackwell (1965) to rewrite the Markov decision process $\langle \mathcal{X}, \mathcal{A}', P, r, \delta \rangle$ we can map our action space $\mathcal{A}'$ to one containing only two actions (which correspond to clearing and waiting). Thus, the results of Blackwell apply and an optimal policy π^* must exist. Let π^* denote any optimal policy of the Markov decision process $\langle \mathcal{X}, \mathcal{A}', P, r, \delta \rangle$. The optimal policy is stationary (the arrival process is stationary and the discount factor is constant) and must clear the market whenever it is in a state of the form $(x_E, 0)$, with $x_E \in \mathbb{N}$. 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

$$\begin{aligned} V_{\pi^*}^U(\boldsymbol{x}^*) = \delta \big[& p^2(r(\boldsymbol{x}^*) + 1 + V_{\pi^*}^U(\mathbf{0})) + p(1-p)(r(\boldsymbol{x}^*) + \Delta_\alpha + V_{\pi^*}^U(\mathbf{0})) \\ & + p(1-p)(r(\boldsymbol{x}^*) + 1 - \Delta_\alpha + V_{\pi^*}^U(\mathbf{0})) + (1-p)^2 V_{\pi^*}^U(\boldsymbol{x}^*) \big] . \end{aligned} \tag{22}$$

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,

$$V_{\pi^*}^U(\boldsymbol{x}^*) \geq r(\boldsymbol{x}^*) + V_{\pi^*}^U(\mathbf{0}). \tag{23}$$

Notice that the right-hand sides of (22) and (23) depend directly on $\boldsymbol{x}^*$ only through $r(\boldsymbol{x}^*)$. Replace $r(\boldsymbol{x}^*)$ with τ^* in (22) and (23) and suppose (23) holds with equality. Then, for every cutoff state $\boldsymbol{x}^*$, $r(\boldsymbol{x}^*) \leq \tau^*$. Using the definition of τ^*, substituting (23) into (22) and rearranging, it can be shown that τ^* satisfies

$$\tau^* + V_{\pi^*}^U(\mathbf{0}) = \frac{\delta p}{1-\delta}. \tag{24}$$

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 (24).

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 $Q_{\tau_j^{\hat{h}}}^{B_i}(\hat{h}_{\tau_j^{\hat{h}}}) = 1$ 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 $k \in \mathbb{N}$. Similarly 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). □

A.11 Proof of Corollary 4

Proof. Under fixed frequency market clearing threshold policies are trivially optimal. We can repeat the procedure from the proof of Theorem 2 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. □

A.12 Proof of Theorem 3

A more precise but not necessarily more transparent statement of the theorem is: *If $\Delta_1 > 0$, there exist $\delta_1, \delta_2 \in (0,1)$ with $\delta_1 < \delta_2$ such that (i) $W^{D,1}(\delta) = W^{0,0}(\delta)$ for $\delta \in [0,\delta_1] \cup \{\delta_2\}$, (ii) $W^{D,1}(\delta) < W^{0,0}(\delta)$ for $\delta \in (\delta_1,\delta_2)$ and (iii) $W^{D,1}(\delta) > W^{0,0}(\delta)$ for $\delta \in (\delta_2, 1]$. If $\Delta_1 \leq 0$ there exists $\delta_3 \in (0,1)$ such that (iv) $W^{D,1}(\delta) < W^{0,0}(\delta)$ for $\delta \in [0,\delta_3)$, (v) $W^{D,1}(\delta_3) = W^{0,0}(\delta_3)$ and (vi) $W^{D,1}(\delta_3) > W^{0,0}(\delta_3)$ for $\delta \in (\delta_3, 1]$.*

We proceed by proving this statement of the theorem, which implies the version as stated in Theorem 3. However, we first state and prove two lemmas.

Lemma 1. *Under discriminatory market clearing, for any threshold policy with threshold τ,*

$$\frac{dV_\tau^D(\tau)}{d\delta} > \cdots > \frac{dV_\tau^D(1)}{d\delta} > \cdots > \frac{dV_\tau^D(0)}{d\delta}.$$

Note that this lemma is well-defined since the value function is differentiable (see, for example Stokey et al. (1989)).

Proof. Let $\boldsymbol{V}_\tau^D$ denote the transpose of $(V_\tau^D(0), V_\tau^D(1), \ldots, V_\tau^D(\tau))$ and $\boldsymbol{r}$ denote the transpose of $(p^2, p, \ldots, p, p(1+(1-p)\Delta_\alpha))$, which is a vector of length $\tau+1$. Furthermore, let $\boldsymbol{I}$ denote the identity matrix. Using $\boldsymbol{P}$ we can rewrite (12) in matrix form,

$$(\boldsymbol{I} - \delta\boldsymbol{P})\boldsymbol{V}_\tau^D = \delta\boldsymbol{r}. \tag{25}$$

This is a special case of Poisson's equation for discounted discrete infinite horizon Markov decision processes.[41]

[41] This may be rewritten as the standard discrete-time Poisson equation (see, for example, Glynn and

Recall that for any threshold policy τ, $\boldsymbol{V}_\tau^D$ satisfies the Poisson equation given by

$$(\boldsymbol{I} - \delta \boldsymbol{P})\boldsymbol{V}_\tau^D = \delta \boldsymbol{r}.$$

Here, $\boldsymbol{V}_\tau^D$ is such that, $i, j \in \{0, 1, \ldots, \tau\}$ with $i > j$, $V_\tau^D(i) > V_\tau^D(j)$. This immediately follows from the fact that $r_i \geq r_j$ (the expected immediate reward earned in the next period is weakly larger for state i compared to state j) and the structure of $\boldsymbol{P}$ (starting from state i, the market transitions to a higher state in expectation compare to state j).

Differentiating the Poisson equation with respect to δ, we obtain

$$(\boldsymbol{I} - \delta \boldsymbol{P})\frac{d\boldsymbol{V}_\tau^D}{d\delta} = \boldsymbol{r} + \boldsymbol{P}\boldsymbol{V}_\tau^D = \frac{\boldsymbol{V}_\tau^D}{\delta}. \tag{26}$$

Therefore, aside from different right-hand-side vectors, $d\boldsymbol{V}_\tau^D/d\delta$ and $\boldsymbol{V}_\tau^D$ satisfy the same Poisson equation.[42] Since $V_\tau^D(i) > V_\tau^D(j)$ we must have that $dV_\tau^D(i)/d\delta > dV_\tau^D(j)/d\delta$. □

Before stating the next lemma, we introduce the following notation. Let $\tau^{D,0}(\delta)$ and $\tau^{D,1}(\delta)$ denote the optimal thresholds under welfare-maximizing and profit-maximizing discriminatory market clearing, respectively.

Lemma 2. *For any $\hat{\delta} \in [0, 1]$ such that $W^{D,1}(\hat{\delta}) \geq W^{0,0}(\hat{\delta})$ and $\tau^{D,1}(\hat{\delta}) > 0$,*

$$\left.\frac{dW^{D,1}(\delta)}{d\delta}\right|_{\delta=\hat{\delta}} > \left.\frac{dW^{0,0}(\delta)}{d\delta}\right|_{\delta=\hat{\delta}}.$$

Proof. Suppose that $\hat{\delta}$ is such that $W^{D,1}(\hat{\delta}) \geq W^{0,0}(\hat{\delta})$ and $\tau^{D,1}(\hat{\delta}) > 0$. We consider (13) with $\tau = \tau^{D,1}(\hat{\delta})$. Differentiating (13) gives

$$\frac{dV_\tau^D(0)}{d\delta} = \frac{V_\tau^D(0)}{\delta(1 - \delta(1 - 2p(1-p)))} + \frac{2\delta p(1-p)}{1 - \delta(1 - 2p(1-p))}\frac{dV_\tau^D(1)}{d\delta}. \tag{27}$$

Since $V_\tau^D(1) > V_\tau^D(0)$ by Lemma 1, (27) implies

$$\frac{dV_\tau^D(0)}{d\delta} > \frac{V_\tau^D(0)}{\delta(1-\delta)}. \tag{28}$$

Meyn, 1996) by adding to the Markov chain an absorbing state in which a reward of 0 is earned. We then suppose that in every period the Markov chain transitions to this absorbing state with probability $1 - \delta$ and scale all other transition probabilities by δ.

[42]Equivalently, $d\boldsymbol{V}_\tau^D/d\delta$ is the vector of value functions corresponding to Markov decision process which, aside from the reward function, is identical to $\langle \mathcal{X}, \mathcal{A}, P, r, \delta \rangle$.

At the point $\hat{\delta}$ we have $\hat{\delta} W^{D,1}(\hat{\delta}) = V_\tau^D(0)\big|_{\hat{\delta}} = \hat{\delta} W^{0,0}(\hat{\delta})$. Furthermore,

$$\delta W^{0,0}(\delta) = \frac{\delta(p^2 + 2p(1-p)\Delta_0)}{(1-\delta)} \Rightarrow \frac{d(\delta W^{0,0}(\delta))}{d\delta} = \frac{p^2 + 2p(1-p)\Delta_0}{(1-\delta)^2}.$$

Combining this with (28) we obtain

$$\left.\frac{d\left(\delta W^{D,1}(\delta)\right)}{d\delta}\right|_{\delta=\hat{\delta}} > \frac{W^{0,0}(\hat{\delta})}{1-\hat{\delta}} = \left.\frac{d\left(\delta W^{0,0}(\delta)\right)}{d\delta}\right|_{\delta=\hat{\delta}}.$$

This implies

$$\left.\frac{dW^{D,1}(\delta)}{d\delta}\right|_{\delta=\hat{\delta}} > \left.\frac{dW^{0,0}(\delta)}{d\delta}\right|_{\delta=\hat{\delta}}$$

as required. □

We now prove Theorem 3. Recall that we let $\tau^{D,0}(\delta)$ and $\tau^{D,1}(\delta)$ denote the optimal thresholds under welfare-maximizing and profit-maximizing discriminatory market clearing, respectively.

Proof. We first deal with the $\Delta_1 > 0$ case. The functions $W^{D,0}$, $W^{D,1}$ and $W^{0,0}$ are continuous (see, for example, Stokey et al. (1989)) and satisfy $W^{D,0}(0) = W^{D,1}(0) = W^{0,0}(0)$. Furthermore, $W^{D,0}$ and $W^{0,0}$ are increasing in δ. By Proposition 3, $\tau^{D,0}(\delta)$ and $\tau^{D,1}(\delta)$ are increasing in δ and by Corollary 1, $\tau^{D,1}(\delta) \geq \tau^{D,0}(\delta)$. Thus, setting $\underline{\delta} = \sup\{\delta \in [0,1] : \tau^{D,1}(\delta) = 0\}$ we have, for sufficiently small $\epsilon > 0$, $\tau^{D,0}(\underline{\delta} + \epsilon) = 0$ and $W^{0,0}(\underline{\delta} + \epsilon) = W^{D,0}(\underline{\delta} + \epsilon) > W^{D,1}(\underline{\delta} + \epsilon)$. Next, we set $\overline{\delta} = \min\{\delta \in (\underline{\delta}, 1] : W^{D,1}(\delta) = W^{0,0}(\delta)\}$ so that $W^{0,0}(\delta) > W^{D,1}(\delta)$ for $\delta \in (\underline{\delta}, \overline{\delta})$. To see that δ_2 exists, notice that under decentralized trade expected welfare per period is constant and equal to $p^2 + 2p(1-p)\Delta_0$. As $\delta \to 1$ expected welfare per period under the profit-maximization mechanism converges to p. Since $p > p^2 + 2p(1-p)\Delta_1$, there exists $\hat{\delta} \in (\delta, 1]$ such that $W^{D,1}(\delta) = W^{0,0}(\delta)$. That $W^{D,1}(\delta) > W^{0,0}(\delta)$ for $\delta > \overline{\delta}$ follows directly from Lemma 2. If $\Delta_1 \leq 0$ we may set $\underline{\delta} = 0$ and the existence of $\overline{\delta}$ follows directly from the previous argument, as well as the fact that in this case $W^{D,0}(0) > W^{D,1}(0)$. □

A.13 Proof of Theorem 4

We start by stating and proving a useful lemma.

Lemma 3. *Let $V^D(0)$ and $V^U(\mathbf{0})$ denote the expected discounted reward earned, starting from an empty market, under the optimal market clearing policy with discriminatory market clearing and uniform market clearing respectively. Then,*

$$\frac{dV^D(0)}{d\delta} > \frac{dV^U(\mathbf{0})}{d\delta}.$$

Note that this lemma is well-defined since the value function is differentiable (see, for example Stokey et al. (1989)).

Proof. By the proof of Lemma 1 we have

$$(\boldsymbol{I} - \delta \boldsymbol{P})\frac{d\boldsymbol{V}^D}{d\delta} = \frac{\boldsymbol{V}^D}{\delta},$$

which shows that, aside from differing reward functions, $d\boldsymbol{V}^D/d\delta$ and $\boldsymbol{V}^D$ solve the same Markov decision process. Similarly, by the analysis in Appendix A.9 we have

$$(\boldsymbol{I} - \delta \boldsymbol{P}')\frac{d\boldsymbol{V}^U}{d\delta} = \frac{\boldsymbol{V}^U}{\delta}$$

and, aside from differing reward functions, $d\boldsymbol{V}^U/d\delta$ and $\boldsymbol{V}^U$ also solve the same Markov decision process. Furthermore, for any $x_S \in \{0, 1, \ldots, \tau^*\}$ we must have $V^D(x_S) > V^U(0, x_S)$ since $\boldsymbol{V}^U$ solves a more constrained version of the same optimization problem as $\boldsymbol{V}^D$. Therefore, $d\boldsymbol{V}^U/d\delta$ solves a more constrained optimization problem for which smaller rewards are earned when the process transitions between states of the associated Markov decision process compared to $d\boldsymbol{V}^D/d\delta$. It immediately follows that $dV^D(0)/d\delta > dV^U(\mathbf{0})/d\delta$. □

We are now ready to prove Theorem 4.

Proof. Since $W^{D,\alpha}(0) = W^{U,\alpha}(0)$, $W^{D,\alpha}(\delta) = V^D(0)/\delta$ and $W^{U,\alpha}(\delta) = V^U(\mathbf{0})/\delta$ the fact that $W^{D,\alpha}(\delta) - W^{U,\alpha}(\delta)$ is positive, increasing in δ and diverges as $\delta \to 1$ follows directly from Lemma 3. We may show that the function $W^{U,\alpha}(\delta) - W^{F,\alpha}(\delta)$ has the required properties by adapting the arguments from the proof of Lemma 3 to uniform and fixed frequency market clearing. By (16) and (17) we have

$$\frac{dW^{F,\alpha}(\delta)}{d\delta} = \frac{\delta^{\tau^*} + \tau^* - 1}{\delta(1-\delta^{\tau^*})}W^{F,\alpha}(\delta) \quad \text{and} \quad \frac{dW^{0,\alpha}(\delta)}{d\delta} = \frac{1}{1-\delta}W^{0,\alpha}(\delta).$$

Since $(\delta^{\tau^*} + \tau^* - 1)/\delta(1-\delta^{\tau^*}) \geq 1/(1-\delta)$ and $W^{F,\alpha}(\delta) \geq W^{D,\alpha}(\delta)$ for all $\delta \in (0,1)$, it follows immediately that $W^{F,\alpha}(\delta) - W^{0,\alpha}(\delta)$ is positive and increasing in δ. The last statement of the theorem follows immediately from the fact that under discriminatory, uniform and fixed frequency market clearing per period welfare must converge to p as $\delta \to 1$ and under discriminatory market clearing per period welfare must converge to $p^2 + 2p(1-p)\Delta_0$. This also shows that $W^{F,\alpha}(\delta) - W^{0,\alpha}(\delta)$ diverges as $\delta \to 1$. □

A.14 Proof of Corollary 5

Proof. For sufficiently δ, $W^{D,0}(\delta) > W^{U,0}(\delta)$. Further, $\lim_{\delta \to 1} \left(W^{D,0}(\delta) - W^{D,1}(\delta) \right) = 0$. Thus, by the results of Theorem 4, for sufficiently large δ, $W^{D,1}(\delta) > W^{U,0}(\delta)$. The other cases are analogous. □

A.15 Proof of Proposition 6

Proof. First, note that for $\delta = 1$ the mechanism runs a surplus in every period since we have discrete type spaces with a continuum of traders. Second, the level of expected discounted welfare generated by the mechanism must increase in δ. To see this, notice that if we hold the allocation rule fixed, an increase in δ increases expected discounted social welfare. If we allow the allocation rule to adjust to the optimal allocation rule following the increase in δ, this will serve only to increase welfare further. Thus, expected discounted social welfare is increasing in δ. It then follows immediately from regularity that the level of expected discounted profit must also increase in δ. □

B Algorithms

B.1 Discriminatory Market Clearing

We begin by describing an algorithm which may be used to compute the optimal threshold τ^* under discriminatory market clearing.

Algorithm 1. *Begin with the threshold policy characterized by* $\tau = 1$ *and solve the linear system defined in* (12). *If* $V_1^D(1) > \Delta_\alpha + V_1^D(0)$, *proceed to step 2. Otherwise, return* $\tau^* = 0$. *At step i,*

1. *Solve* (12) *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 1 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.

B.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 $V_\tau^U(y_E, y_S)$ 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

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

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,

$$V_\tau^U(\boldsymbol{y}) = \delta \sum_{\boldsymbol{z} \in Z} P_Z(\boldsymbol{z}) \left[V_\tau^U(\boldsymbol{y}+\boldsymbol{z}) \mathbb{1}(r(\boldsymbol{y}+\boldsymbol{z}) < \tau) \right.$$
$$\left. + (r(\boldsymbol{y}+\boldsymbol{z}) + V_\tau^U(\mathbf{0})) \mathbb{1}(r(\boldsymbol{y}+\boldsymbol{z}) \geq \tau) \right]. \tag{30}$$

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

$$V_{\tau^*}^U(\boldsymbol{x}^*) > r(\boldsymbol{x}^*) + V_{\tau^*}^U(\mathbf{0}) \quad \text{and} \quad V_{\tau^*}^U(x_E^*, x_S^*+1) \leq r(x_E^*, x_S^*+1) + V_{\tau^*}^U(\mathbf{0}).$$

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 of non-identical suboptimal pair arrives. Since τ^* applies to all cutoff states, to compute τ^* it suffices to find a single cutoff state. Algorithm 2 determines τ^* by computing the cutoff state $(0, x_S^*)$ using the aforementioned stopping condition.

Algorithm 2. *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 2 may not be the most economical algorithm. For example, for larger value of δ, a more computationally 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.

B.3 Fixed Frequency Market Clearing

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

Algorithm 3. *Begin with the threshold policy characterized by $\tau = 2$ and compute $W^F(2)$ and $W^F(1)$ using* (16)*. 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* (16)*.*
2. *If $W^F(i) \geq W^F(i-1)$, proceed to step $i+1$. Otherwise, return $\tau^* = i-1$.*

C Further Extensions

C.1 Renewal Process Arrivals

The results of this paper immediately generalize to the case in which pairs of buyers and sellers arrive according to a Poisson process.[43] If the intensity of the arrival process is η then the expected inter-arrival time is $1/\eta$ and the results in Section 4.1 apply 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, at any time s such that $\boldsymbol{E}[A(s)] > \boldsymbol{E}[A(0)]$, the

[43]In Appendix D we further generalize this to the case in which buyers and sellers arrive according to independent Poisson processes.

market designer must recompute the optimal policy by applying a discount factor of $\delta^{\boldsymbol{E}[A(s)]}$ to the next arrival (but retaining a discount factor of $\delta^{\boldsymbol{E}[A(0)]}$ for all subsequent arrivals).

C.2 Group Arrivals

By appropriately updating the transition probabilities of the underlying Markov decision processes, our baseline model can easily be extended to allow for a group of pairs to arrive in each period. Let N denote the maximum number of pairs who arrive in a given period and $\xi = (\xi_E, \xi_S)$ denote the number of efficient and suboptimal pairs who arrive in a given period, where $\xi \in \mathcal{Z} := \{(\xi_E, \xi_S) \in \mathbb{Z}^2_{\geq 0} : \xi_E + \xi_S \leq N\}$. Under discriminatory market clearing, the optimal policy will be a threshold policy in which up to τ identical suboptimal pairs are stored. For $x_S \in \{0, \ldots, \tau\}$, the corresponding threshold policy is characterized by the linear system

$$\begin{aligned} V^D_\tau(x_S) = \delta \sum_{\xi \in \mathcal{Z}} P(\Xi = \xi) \big[&(\xi_E + V^D_\tau(x_S + \xi_S))\mathbb{1}(x_S + \xi_S \leq \tau) \\ &+ (\xi_E + \Delta_\alpha(x_S + \xi_S - \tau) + V^D_\tau(\tau))\mathbb{1}(x_S + \xi_S > \tau)\big]. \end{aligned}$$

Similarly, under uniform market clearing, the optimal policy will be a threshold policy in which trades accumulate in the order book up to a threshold value τ. For $(x_E, x_S) \in \mathbb{Z}^2_{\geq 0}$ be such that $x_E + \Delta_\alpha x_S < \tau$ and $x_S > 0$, the corresponding threshold policy is characterized by

$$\begin{aligned} V^U_\tau(x_E, x_S) = \delta \sum_{\xi \in \mathcal{Z}} P(\Xi = \xi) \big[&V^U_\tau(x_E + \xi_E, x_S + \xi_S)\mathbb{1}(x_E + \xi_E + \Delta_\alpha(x_S + \xi_S) < T) \\ &+ (x_E + \xi_E + \Delta_\alpha(x_S + \xi_S) + V^U_\tau(0,0))\mathbb{1}(x_E + \xi_E + \Delta_\alpha(x_S + \xi_S) \geq T)\big]. \end{aligned}$$

As before, fixed frequency market clearing can be dealt with by appropriately modifying (16). Considering group arrivals essentially enables us to consider non-uniform arrival processes. For example, this extension be used to model markets in which large numbers of buyers and sellers tend to arrive together.[44]

[44] 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.

C.3 Multi-Unit Traders

With binary valuations and costs, it is also 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 $\text{supp}(G) = \{0, \ldots, k\}$.[45]

The arrival of the period t buyer and seller is 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.

C.4 Demand Complementarity

When buyer demand exhibits complementaries, this increases the relative benefit of storing traders in dynamic environments. Assume that a certain proportion of buyers demand two units of the good and derive zero utility from consuming a single unit. This increases the benefit of delaying market clearing since there are additional gains from rematching such traders. Specifically, assume that β is the proportion of buyers who demand two units with constant marginal value. Then when a given buyer arrives, with probability βp demand is $(\overline{v}, \overline{v})$, with probability $\beta(1-p)$ demand is $(\underline{v}, \underline{v})$, with probability $(1-\beta)p$ demand is $\overline{v}$ and with probability $(1-\beta)(1-p)$ demand is $\underline{v}$. We assume the type distribution for sellers is unchanged.

When $\delta = 0$, buyers who demand two units cannot trade and the designer's payoff is $(1-\beta)(p^2 + 2p(1-p)\Delta_\alpha)$. When $\delta = 1$, there is demand for $(1+\beta)p$ units at valuation $\overline{v}$ so under the efficient allocation p sellers of type $\underline{c}$ and βp sellers of type $\overline{c}$ trade. For $0 < \delta < 1$,

[45] 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 $\text{Bn}(k, p)$.

we consider discriminatory market clearing and suppose the seller side of the market is short.[46] After an arrival, the designer should clear any available $(\overline{v}, \overline{v}, \underline{c}, \underline{c})$ matchings, then clear any $(\overline{v}, \underline{c})$ pairs and store $(\overline{v}, \overline{v}, \underline{c}, \overline{c})$, $(\overline{v}, \overline{v}, \overline{c}, \overline{c})$ and $(\overline{v}, \overline{c})$ matches up to a threshold. The optimal threshold for each type of match decreases in the present value of that match. If a $(\underline{v}, \overline{c})$ pair arrives, the seller should always be stored.[47] Finally, notice that matches of type $(\overline{v}, \overline{v}, \underline{c})$ and $(\overline{v}, \overline{v}, \overline{c})$ can be stored at no cost and there is no bound on the number of matches of this type that can accumulate under the optimal mechanism. Thus, in the case of uniform market clearing the details of deriving the optimal mechanism are similar to the case in which unpaired agents arrive at the market (see Appendix D).

D Methodology for Dealing with Unpaired Agents

To deal with the problem of unpaired agents under uniform market clearing, we identify a finite partition of the state space which can be used to determine the optimal policy. One can accomplish this by computing the optimal threshold policy of a related Markov decision process. One can compute the optimal threshold $\overline{\tau}$ of this Markov decision process using the results of Section 4.1. Agents of type $\overline{v}$ and $\underline{c}$ arrive with equal probability under this new and the original Markov decision processes. However, rematching frictions are reduced under the new Markov decision process because all agents must arrive in pairs. Thus, the threshold $\overline{\tau}$ provides an upper bound on the value of trades which can be stored under the optimal policy of the original Markov decision process. This allows a finite number of candidate optimal policies to be identified so that a simple policy iteration algorithm can be used to determine the optimal policy.

For ease of exposition, we now consider a continuous-time version of the model and suppose that buyers arrive according to the Poisson process $\{N_t\}_{t\in\mathbb{R}_{\geq 0}}$ with rate η and sellers arrive according to the Poisson process $\{M_t\}_{t\in\mathbb{R}_{\geq 0}}$ with rate η. This ensures that a pair of buyers and sellers arrives with probability zero, which reduces the number of transitions which must be considered in the Markov Decision Process. However, the methodology de-

[46]When the buyer side of the market is short, things are simple. We match $\overline{v}$ and $(\overline{v}, \overline{v})$ types with the appropriate number of $\underline{c}$ types until the seller side of the market is short.

[47]Sellers of type $\underline{c}$ are useful since this type of seller will trade with non-zero probability even when $\delta = 1$.

scribed here immediately applies to the setup described in Section 5

We let $\{B_i\}_{i\in\mathbb{N}}$ denote the set of buyers, so that the arrival time of B_i is given by $\min\{t \in \mathbb{R}_{\geq 0} : N_t = i\}$. Similarly, let $\{S_j\}_{j\in\mathbb{N}}$ denote the set of sellers, with the arrival of S_j given by $\min\{t \in \mathbb{R}_{\geq 0} : M_t = j\}$. All other details of the setup remain unchanged.

D.1 Markov Decision Process

We start by defining the state space of the Markov decision process. We can think of the state space as as having two branches, depending on whether there is a shortage of $\overline{v}$ or $\underline{c}$ agents. Since buyers and sellers arrive at the same rate and agents of type $\overline{v}$ and $\underline{c}$ arrive with equal probability, by symmetry, there is no need to distinguish these cases. Denote by $\mathcal{Y}_1 = \{(y_E, y_S, u_E, u_S, u_S') : y_S + u_E \neq 0, u_E u_S = 0\}$ the region of the state space in which there is a shortage of either $\overline{v}$ or $\underline{c}$ agents. If we have a shortage $\underline{c}$ agents, y_E denotes the number of $(\overline{v}, \underline{c})$ pairs, y_S denotes the number of $(\overline{v}, \overline{c})$ pairs and u_E, u_S and u_S' denote the number of unpaired $\overline{v}$, $\overline{c}$ and $\underline{v}$ agents, respectively. If we have a shortage of $\overline{v}$ agents, y_E denotes the number of $(\overline{v}, \underline{c})$ pairs, y_S denotes the number of $(\underline{v}, \underline{c})$ pairs and u_E, u_S and u_S' denote the number of unpaired $\underline{c}$, $\underline{v}$ and $\overline{c}$ agents, respectively.

There are two additional regions of the state space where these branches join. Let $\mathcal{Y}_0 = \{(w_S, w_S') : w_S, w_S' \in \mathbb{Z}_{\geq 0}\}$ denote the region of the state space in which there are no $\overline{v}$ or $\underline{c}$ agents present. Here, w_S counts the number of $\underline{v}$ agents and w_S' counts the number of $\overline{c}$ agents. If the process is in $\mathcal{Y}_0$ and a $\overline{v}$ or $\underline{c}$ arrival occurs, the process will transition to the appropriate branch of $\mathcal{Y}_1$. Finally, the process may transition from $\mathcal{Y}_1$ to a region in which all $\overline{v}$ and $\underline{c}$ agents are part of an efficient pair. We denote this region of the state space by $\mathcal{Y}_2 = \{(y_E, 0, 0, u_S, u_S') : y_E > 0\}$. For convenience, we assume that if a transition from $\mathcal{Y}_1$ to $\mathcal{Y}_2$ occurs, the labeling of the states is unchanged. If the process is in $\mathcal{Y}_2$ and a $\overline{v}$ or $\underline{c}$ arrival takes place, the process will transition to the appropriate branch of $\mathcal{Y}_1$. The state space of the Markov decision process is given by $\mathcal{Y} = \mathcal{Y}_0 \cup \mathcal{Y}_1 \cup \mathcal{Y}_2$. Although this notation seems cumbersome, defining the state space in this manner will be convenient when we partition the state space and compute the optimal policy.

We now define the set of actions available in each state. For states $\boldsymbol{y} \in \mathcal{Y}_0$ we can set $\mathcal{B}_{\boldsymbol{y}} = \{\boldsymbol{0}\}$ since $r(\boldsymbol{y}) = 0$ and it is not optimal for the designer to clear. For states

$\boldsymbol{y} = (y_E, y_S, u_E, u_S, u'_S) \in \mathcal{Y}_1 \cup \mathcal{Y}_2$ we set $\mathcal{B}_{\boldsymbol{y}} = \{\boldsymbol{y}, \mathbf{0}\}$.

Let $P_{\boldsymbol{b}}(\boldsymbol{y}, \boldsymbol{y}')$ denote the probability that if the designer takes action $\boldsymbol{b}$ in state $\boldsymbol{y}$, the state of the market following the next Poisson arrival is $\boldsymbol{y}'$. For $\boldsymbol{y} = (w_S, w'_S) \in \mathcal{Y}_0$, if a $\overline{c}$ or $\underline{v}$ arrival occurs, the process will remain in $\mathcal{Y}_0$ and we have

$$P_{\boldsymbol{b}}(\boldsymbol{y}, (w_S + 1, w'_S)) = (1-p)/2 \quad \text{and} \quad P_{\boldsymbol{b}}(\boldsymbol{y}, (w_S, w'_S + 1)) = (1-p)/2.$$

If a $\overline{v}$ or $\underline{c}$ arrival occurs, the process will transition to $\mathcal{Y}_1$ and we have

$$P_{\boldsymbol{b}}(\boldsymbol{y}, (0, 0, 1, w_S, w'_S)) = p/2 \quad \text{and} \quad P_{\boldsymbol{b}}(\boldsymbol{y}, (0, 0, 1, w'_S, w_S)) = p/2.$$

For $\boldsymbol{y} = (y_E, y_S, u_E, u_S, u'_S) \in \mathcal{Y}_1 \cup \mathcal{Y}_2$ and $\boldsymbol{b} = \boldsymbol{y}$, the process moves to $\mathbf{0}$ and the previous transition probabilities apply. We next consider $\boldsymbol{y} \in \mathcal{Y}_1$ and $\boldsymbol{b} = \mathbf{0}$ with $u_E > 0$. If unpaired agents of type $\overline{v}$ are stored and a $\overline{v}$ or $\underline{v}$ arrival takes place, no new pairs can be created. By symmetry, we have

$$P_{\mathbf{0}}(\boldsymbol{y}, \boldsymbol{y} + (0, 0, 1, 0, 0)) = p/2 \quad \text{and} \quad P_{\mathbf{0}}(\boldsymbol{y}, \boldsymbol{y} + (0, 0, 0, 0, 1)) = (1-p)/2.$$

If unpaired agents of type $\overline{v}$ are stored and a $\underline{c}$ arrival takes place, an efficient pair is created. If a $\overline{c}$ arrival occurs, a suboptimal pair is created. By symmetry, we have

$$P_{\mathbf{0}}(\boldsymbol{y}, \boldsymbol{y} + (1, 0, -1, 0, 0)) = p/2 \quad \text{and} \quad P_{\mathbf{0}}(\boldsymbol{y}, \boldsymbol{y} + (0, 1, -1, 0, 0)) = (1-p)/2.$$

Similar logic applies when $u_S > 0$ and we have

$$P_{\mathbf{0}}(\boldsymbol{y}, \boldsymbol{y} + (0, 1, 0, -1, 0)) = p/2, \quad P_{\mathbf{0}}(\boldsymbol{y}, \boldsymbol{y} + (1, -1, 0, 1, 0)) = p/2,$$
$$P_{\mathbf{0}}(\boldsymbol{y}, \boldsymbol{y} + (0, 0, 0, 1, 0))) = (1-p)/2, \quad P_{\mathbf{0}}(\boldsymbol{y}, \boldsymbol{y} + (0, 0, 0, 0, 1)) = (1-p)/2.$$

Similarly, if $u_E = u_S = 0$ we have

$$P_{\mathbf{0}}(\boldsymbol{y}, \boldsymbol{y} + (0, 0, 1, 0, 0)) = p/2, \quad P_{\mathbf{0}}(\boldsymbol{y}, \boldsymbol{y} + (1, -1, 0, 0, 0)) = p/2,$$
$$P_{\mathbf{0}}(\boldsymbol{y}, \boldsymbol{y} + (0, 0, 0, 1, 0)) = (1-p)/2, \quad P_{\mathbf{0}}(\boldsymbol{y}, \boldsymbol{y} + (0, 0, 0, 0, 1)) = (1-p)/2.$$

If $\boldsymbol{y} \in \mathcal{Y}_2$ and $\boldsymbol{b} = \mathbf{0}$, we have

$$P_{\mathbf{0}}(\boldsymbol{y}, \boldsymbol{y} + (0, 0, 0, 1, 0)) = (1-p)/2 \quad \text{and} \quad P_{\mathbf{0}}(\boldsymbol{y}, \boldsymbol{y} + (0, 0, 0, 0, 1)) = (1-p)/2.$$

If $u_S > 0$ we have $P_{\mathbf{0}}(\boldsymbol{y}, (y_E, 1, 0, u_S - 1, u'_S)) = p/2$. Otherwise, $P_{\mathbf{0}}(\boldsymbol{y}, (y_E, 0, 1, 0, u'_S)) = p/2$. Finally, $P_{\mathbf{0}}(\boldsymbol{y}, (y_E, 1, 0, u'_S - 1, u_S)) = p/2$ if $u'_S > 0$ and $P_{\mathbf{0}}(\boldsymbol{y}, (y_E, 0, 1, 0, u_S)) = p/2$ otherwise.

Finally, if $\boldsymbol{y} \in \mathcal{Y}_1 \cup \mathcal{Y}_2$ and action $\boldsymbol{b} = \boldsymbol{y}$ is implemented, the immediate reward is

$$s(\boldsymbol{y}) = y_E + \Delta_\alpha y_S.$$

Otherwise, the designer earns no reward. The market designer's problem is to determine the optimal policy π^* of the Markov decision process $\langle \mathcal{Y}, \mathcal{B}, P, s, \delta \rangle$. However, dealing with unpaired agents is problematic. From the designer's perspective, no immediate reward is earned when unpaired agents are cleared from the market but they are useful for forming pairs in the future. Therefore, there is no upper bound on the number of unpaired agents stored under the optimal policy.

D.2 Partitioning the State Space

To deal with the problem of unpaired agents, we identify a finite partition of the state space which can be used to determine the optimal policy. We accomplish this by computing the optimal threshold policy of a related Markov decision process.

Suppose temporarily that every buyer arrives paired with a seller of type $\overline{c}$ and every seller arrives paired with a buyer of type $\underline{v}$, so that agents arrive according to the Poisson process $\{\overline{K}_t\}_{t \in \mathbb{R}_{\geq 0}}$ with rate 2η. Any given arrival is $(\overline{v}, \overline{c})$ with probability $p/2$, $(\underline{v}, \underline{c})$ with probability $p/2$ and $(\underline{v}, \overline{c})$ with probability $(1-p)$. The state space, the set of available actions and the reward function of this Markov decision process are the same as the Markov decision process $\langle \mathcal{X}, \mathcal{A}', P, r, \delta \rangle$ considered in Section 4.1. However, we must update the transition probabilities defined. For our new process we have $\overline{P}_{\boldsymbol{a}}(\boldsymbol{x}, \boldsymbol{x} - \boldsymbol{a}) = 1 - p$. If $\boldsymbol{x} = \boldsymbol{a}$ or $x_S = 0$ we also have

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

Otherwise,

$$\overline{P}_{\boldsymbol{a}}(\boldsymbol{x}, (x_E, x_S + 1)) = p/2 \quad \text{and} \quad \overline{P}_{\boldsymbol{a}}(\boldsymbol{x}, (x_E + 1, x_S - 1)) = p/2.$$

Thus, we have a Markov decision process $\langle \mathcal{X}, \mathcal{A}', \overline{P}, r, \delta \rangle$. Repeating the analysis from Section 4.1 allows us to determine the optimal threshold policy of this Markov decision process $\langle \mathcal{X}, \mathcal{A}', \overline{P}, r, \delta \rangle$. Let $\overline{\tau}$ denote the threshold corresponding to the optimal threshold policy.

Theorem 5. *Under the optimal policy* π^* *of* $\langle \mathcal{Y}, \mathcal{B}, P, s, \delta \rangle$, $\pi^*(\boldsymbol{y}) = \boldsymbol{y}$ *for all* $\boldsymbol{y} \in \mathcal{Y}$ *such that* $s(\boldsymbol{y}) > \overline{\tau}$.

Proof. Suppose the market is in state $\boldsymbol{y} \in \mathcal{Y}$, which might have unpaired agents. We compare the optimal market clearing policy under the original Markov decision process $\langle \mathcal{Y}, \mathcal{B}, P, s, \delta \rangle$ to the optimal market clearing policy when the arrival process is $\{\overline{K}_t\}_{t \in \mathbb{R}_{\geq 0}}$. If the market is in state $\boldsymbol{y}$ and the future arrival process is $\{\overline{K}_t\}_{t \in \mathbb{R}_{\geq 0}}$, it must be optimal to clear the market if $s(\boldsymbol{y}) > \overline{\tau}$ because the threshold $\overline{\tau}$ applies regardless of whether unpaired agents are present (this follows directly from the proof of Theorem 1). The reward earned when the market is cleared immediately is independent of the arrival process but the benefit of waiting is greater under the arrival process $\{\overline{K}_t\}_{t \in \mathbb{R}}$. Although agents of type $\overline{v}$ and $\underline{c}$ arrive at the same rate under each arrival process, these agents always arrive in a pair under $\{\overline{K}_t\}_{t \in \mathbb{R}}$ which improves future trading possibilities. It follows that $\pi^*(\boldsymbol{y}) = \boldsymbol{y}$ for all $\boldsymbol{y} \in \mathcal{Y}$ such that $s(\boldsymbol{y}) > \overline{\tau}$. □

We have already established that it is optimal to wait for all states $\boldsymbol{y} \in \mathcal{Y}_0$ and it is optimal to clear for states $\boldsymbol{y} \in \mathcal{Y}_1 \cup \mathcal{Y}_2$ such that $s(\boldsymbol{y}) > \overline{\tau}$. We now define the finite partition of the state space $\mathcal{Y}$ which contains states for which we must determine whether it is optimal to wait to clear the market. That is, we use the upper bound $\overline{\tau}$ to determine the finite partition over which we must search for the optimal policy.

Suppose the market is in some state $\boldsymbol{y} \in \mathcal{Y}_1 \cup \mathcal{Y}_2$ such that $r(\boldsymbol{y}) \leq \overline{\tau}$. Notice that under the optimal policy π^*

$$\nu_{y_E} = 1 + \max\{i \in \mathbb{N} : y_E + \Delta_\alpha i \leq \overline{\tau}\} \tag{31}$$

is the maximum number of unpaired agents that could feasibly trade when the next market clearing event occurs. It follows that for $i, j \geq \nu_{y_E}$, $V_{\pi^*}(y_E, y_S, i, u_S, u'_S) = V_{\pi^*}(y_E, y_S, j, u_S, u'_S)$. This implies $\pi^*(y_E, y_S, i, u_S, u'_S) = \pi^*(y_E, y_S, j, u_S, u'_S)$. Analogous arguments apply to the u_S and u'_S states. Intuitively, once ν_{y_E} unpaired agents have accumulated, any additional

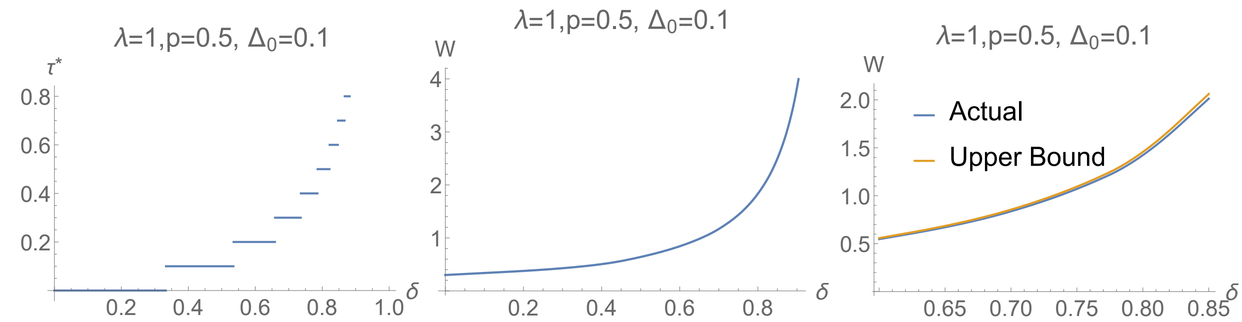

(a) Optimal threshold policy (b) Coupled arrival processes (c) Indep. arrival processes

Figure 12: In Panel (a), the methodology developed in Section 4.1 was used to compute the optimal threshold policy for the continuous-time version of the model with coupled Poisson arrival processes. Panel (b) shows welfare under the optimal threshold policy. Panel (c) shows welfare under the optimal policy when the arrival processes are independent, where the optimal policy was computed using Algorithm 4. This panel also shows welfare under the optimal policy of $\langle \mathcal{X}, \mathcal{A}', \overline{P}, r, \delta \rangle$ (the process used to compute the upper bound $\overline{\tau}$).

unpaired agents are of no benefit to the designer. Therefore, the set of states in $\mathcal{Y}_1$ that must be checked in order to determine the optimal policy is given by

$$\mathcal{Z}_1 = \bigcup_{y_E=0}^{\lfloor \overline{\tau} \rfloor} \bigcup_{y_S=0}^{\nu_{y_E}-1} \bigcup_{i=0}^{\nu_{y_E}-y_S} \left(\left(\bigcup_{j=0}^{\nu_{y_E+y_S+i}} (y_E, y_S, i, 0, j) \right) \bigcup \left(\bigcup_{j=0}^{\nu_{y_E+y_S}} (y_E, y_S, 0, i, j) \right) \right). \tag{32}$$

The set of states in $\mathcal{Y}_2$ that must be checked is given by

$$\mathcal{Z}_2 = \bigcup_{y_E=1}^{\lfloor \overline{\tau} \rfloor} \bigcup_{i=0}^{\nu_{y_E}} \bigcup_{j=0}^{\nu_{y_E}} (y_E, 0, 0, i, j). \tag{33}$$

Theorem 6. *The optimal policy π^* is uniquely determined by its values on the finite subsets $\mathcal{Z}_1 \subset \mathcal{Y}_1$ and $\mathcal{Z}_2 \subset \mathcal{Y}_2$ defined in* (32) *and* (33)*.*

Proof. We have established that for $\boldsymbol{y} \in \mathcal{Y}$ such that $s(\boldsymbol{y}) > \overline{\tau}$, $\pi^*(\boldsymbol{y}) = \boldsymbol{y}$. For states $\boldsymbol{y} \in \mathcal{Y}_0$, we have $\pi^*(\boldsymbol{y}) = \boldsymbol{0}$. Suppose $\pi^*(\boldsymbol{y})$ is known for all $\boldsymbol{y} \in \mathcal{Z}_2$. Then for any $\boldsymbol{y} \in \mathcal{Y}_2$,

$$V_{\pi^*}(\boldsymbol{y}) = V_{\pi^*}(\boldsymbol{y}') \quad \text{and} \quad \pi^*(\boldsymbol{y}) = \pi^*(\boldsymbol{y}'),$$

where $\boldsymbol{y}' = (y_E, 0, 0, \min\{u_S, \nu_{y_E}\}, \min\{u'_S, \nu_{y_E}\}) \in \mathcal{Z}_2$. Analogous logic applies for states $\boldsymbol{y} \in \mathcal{Y}_1$. Thus, π^* is completely characterized. □

Since $\mathcal{Z}$ contains a finite number of states we have a finite number of candidate optimal policies for the Markov decision process $\langle \mathcal{Y}, \mathcal{B}, P, s, \delta \rangle$. We can therefore use a policy iteration algorithm to determine the optimal policy π^*. We start by setting $\mathcal{Z}_0 = \{(i, j, 0) \in \mathcal{Y}_0 : 0 \leq i, j \leq \nu_0\}$, where $\nu_0 = 1 + \max\{i \in \mathbb{N} : \Delta_\alpha i \leq \overline{\tau}\}$

Algorithm 4. *Begin with the policy $\pi_0(\boldsymbol{y}) = \boldsymbol{0}\ \forall\ \boldsymbol{y} \in \mathcal{Z}_0 \cup \mathcal{Z}$. At step $i \in \mathbb{N}$,*

1. *For every $\boldsymbol{y} \in \mathcal{Z}_0 \cup \mathcal{Z}$ compute $V_{\pi_{i-1}}(\boldsymbol{y})$.*
2. *For every $\boldsymbol{y} \in \mathcal{Z}$, if $V_{\pi_{i-1}}(\boldsymbol{y}) \leq r(\boldsymbol{y}) + V_{\pi_{i-1}}(\boldsymbol{0})$ set $\pi_i(\boldsymbol{y}) = \boldsymbol{y}$. Otherwise set $\pi_i(\boldsymbol{y}) = \boldsymbol{0}$.*
3. *If, for every $\boldsymbol{y} \in \mathcal{Z}$, $\pi_i(\boldsymbol{y}) = \pi_{i-1}(\boldsymbol{y})$, return $\pi^* = \pi_i$. Otherwise proceed to step $i + 1$.*

Since $\mathcal{Z}$ is finite, by Howard (1960), this algorithm returns π^* in finitely many steps. For states in $\mathcal{Y}_0$ it suffices to compute $V_\pi(\boldsymbol{y})$ for $\boldsymbol{y} \in \mathcal{Z}_0$. To see this, notice that $V_\pi(w_S, w'_S, 0) = V_\pi(\min\{w_S, \nu_0\}, \min\{w'_S, \nu_0\}, 0)$ since at most ν_0 agents of type $\underline{v}$ or $\overline{c}$ can trade when the next market clearing event occurs. Similarly, for states in $\mathcal{Y}_1 \cup \mathcal{Y}_2$ it suffices to compute values for states in $\mathcal{Z}$.

To compute the value function V_π for a given policy π, we may formulate and solve a finite linear system similar (this was done for $\langle \mathcal{X}, \mathcal{A}, P, r, \delta \rangle$ in (29) and (30)). Alternatively, we can compute V_{π_i} numerically and recursively as follows. Start by setting $V_{\pi,0}(\boldsymbol{y}) = 1$ for all $y \in \mathcal{Z}_0 \cup \mathcal{Z}$. At step j, we set

$$V_{\pi_i,j+1}(\boldsymbol{y}) = \sum_{\boldsymbol{y}' \in \mathcal{Z}_0 \cup \mathcal{Z}} P_{\pi_i(\boldsymbol{y}')}(\boldsymbol{y}, \boldsymbol{y}')\left(s(\boldsymbol{y}') + \delta V_{\pi_i,j}(\boldsymbol{y}')\right)$$

until $|V_{\pi_i,j+1}(\boldsymbol{y}) - V_{\pi_i,j}(\boldsymbol{y})| < \varepsilon$, where ε is the prescribed tolerance.

Some numerical results produced by implementing Algorithm 4 in Matlab R2016a can be found in Figure 12.