Price Discovery in Waiting Lists

Itai Ashlagi* Jacob D. Leshno† Pengyu Qian‡ Amin Saberi§

August 23, 2022¶

Abstract

Waiting lists offer agents a choice between types of items with associated waiting times. These waiting times function as prices and are endogenously determined by a tâtonnement-like price discovery process: an item's price increases when an agent selects it, and decreases when an item arrives and is assigned. We show that this simple price discovery process generates waiting times that fluctuate around market-clearing prices, and that the loss from price fluctuations is bounded by the size of price adjustments. The technical approach and intuition for the results relies on a connection between price adjustments in the waiting list and the stochastic gradient descent optimization algorithm. We further show that this simple price discovery process is asymptotically optimal if the size of price adjustments optimally balances between the adaptivity and the rigidity of the price discovery process.

*Stanford University, iashlagi@stanford.edu.
†University of Chicago, Booth, Jacob.Leshno@ChicagoBooth.edu.
‡Purdue University, pqian20@gmail.com.
§Stanford University, saberi@stanford.edu.

¶This research was initiated during the authors' visit to the Simons Institute for the Theory of Computing as part of the Online and Matching-Based Market Design program. A one-page abstract of an earlier version of this work was published in EC '20 (Proceedings of the 21st ACM Conference on Economics and Computation) under the title "Queue Lengths as Constantly Adapting Prices: Allocative Efficiency Under Random Dynamics". We thank Eric Budish, Emir Kamenica, Yash Kanoria, Chad Syverson, Dan Waldinger, Adam Wierman, and seminar participants at Brown, CalTech, Chicago Booth, Cornell, the Israel Algorithmic Game Theory seminar , Market Design NBER, University of Chicago and Northwestern University Theory conference, PSE workshop on Dynamic Matching and Queueing, University of Tokyo, Northwestern University, NYU Math+Econ+Code seminar, and Warwick Theory workshop for helpful comments and advice. Anand Shah provided invaluable research support. This work is supported by the Robert H. Topel Faculty Research Fund at the University of Chicago Booth School of Business.

1 Introduction

Waiting lists are commonly used to allocate scarce items that arrive over time. To facilitate an efficient assignment of items to agents, waiting lists offer agents a choice between items with associated waiting times. The required waiting cost for an item can be interpreted as the item's price, as waiting times serve a similar role as monetary prices in guiding the allocation and rationing items. For example, the New York City Housing Authority (NYCHA) holds a separate queue for each public housing development, and applicants can join at most one development's queue.[1] Allowing agents to choose developments allows NYCHA to better match applicants to developments, and the waiting list rations desirable projects by requiring higher waiting times for over-demanded projects.

This paper considers waiting times as prices, and studies their endogenous determination under the waiting list. In contrast to monetary prices that can be set by the planner, the required waiting time for an item is endogenously determined by the number of agents queueing for the item. The planner is not required to have any knowledge of market-clearing prices or agent preferences. Instead, the waiting list determines waiting times through an adjustment process that resembles a tâtonnement process: the waiting time for an item increases when an agent joins its queue, and decreases when an item is supplied to an agent that leaves the queue. We view this adjustment process as a price discovery process and characterize the resulting prices and allocation.

To illustrate, consider a simple example.

Example 1. *A single kind of item is allocated by queue. Items arrive according to a Poisson process with rate* 1*. Agents arrive according to a Poisson process with rate* 2*. Agents have heterogeneous values for the item distributed* $v \sim U[0, 1]$ *and a cost of waiting* $c(w) = 0.02 \cdot w$*. Arriving agents choose whether to join the queue or to leave unassigned after observing the length of the queue.*

Intuitively, only half the agents in Example 1 can be assigned items. Items will be assigned to the highest-value agents if all agents face the market-clearing price of 1/2, which is equal to the cost of waiting in the queue when the queue length is[2] 24.

The waiting list has no a priori knowledge of agent preferences or the market-clearing price and, in particular, the queue length is not specified by the planner. Instead, the queue length naturally adjusts over time according to agent decisions. Because the length of the queue continuously changes as agents and items arrive, the price (waiting cost) does

[1] See NYCHA's Tenant Selection and Assignment Plan section VII. part D, available at https://www1.nyc.gov/assets/nycha/downloads/pdf/TSAPlan.pdf (retrieved July 2022).

[2] An agent who observes a queue length of 24 will wait until 25 copies of the item arrive, as the first 24 copies will be assigned to the agents ahead of him. Because each arrival takes one time unit in expectation, the expected waiting cost is $0.02 \cdot 25 = 1/2$.

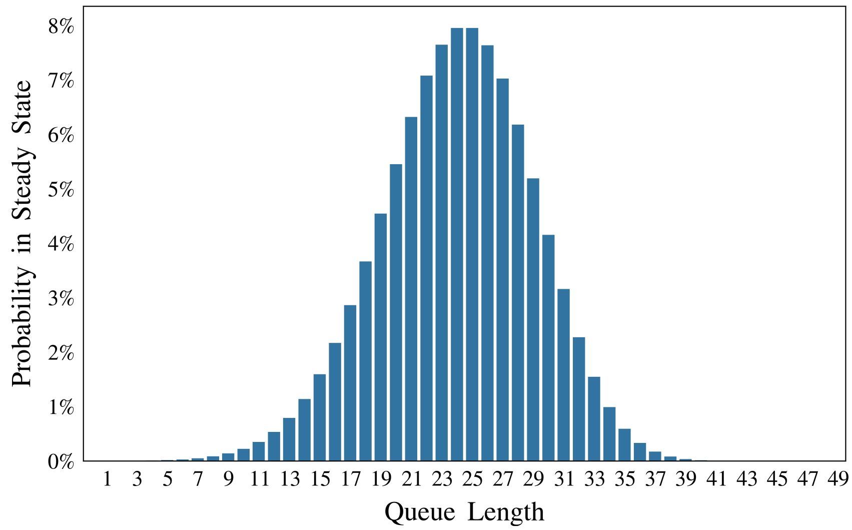

Figure 1: Distribution of queue lengths observed by a randomly drawn agent arriving to the waiting list in Example 1.

not converge and continuously fluctuates over time. Figure 1 presents the probability that an arriving agent observes any possible queue length in steady state. It shows that queue lengths are likely to be close to 24, but fluctuate considerably. These fluctuations cause misallocations and reduce allocative efficiency. Figure 2 shows the implied assignment under the waiting list, and how it differs from the allocative efficiency-maximizing assignment. For example, agents with a value of 0.4 have an 18% probability of arriving when the queue is randomly short; thus they face a price that is lower than 0.4 and are assigned an item. Meanwhile, agents with a value of 0.6 have a 19% probability of arriving when the queue is randomly long; thus they face a price that is higher than 0.6 and are not assigned an item.

To study the waiting list's price discovery process in general economies, we study a model that allows for any number of item types and general agent preferences. In the model agents and items arrive over time. There are finitely many types of items, and the waiting list maintains a queue for each item type. Agents have unit demand and their values for items are drawn i.i.d. from an arbitrary joint value distribution. An arriving agent observes the queue lengths, and can choose to join a single queue to wait until being assigned the respective item.[3]

In general, the waiting list's price discovery process generates waiting times that fluc-

[3] Agents cannot switch queues or otherwise change their choice at a later time. An arriving agent chooses the queue that maximizes her item value minus the expected waiting cost given the current length of its queue.

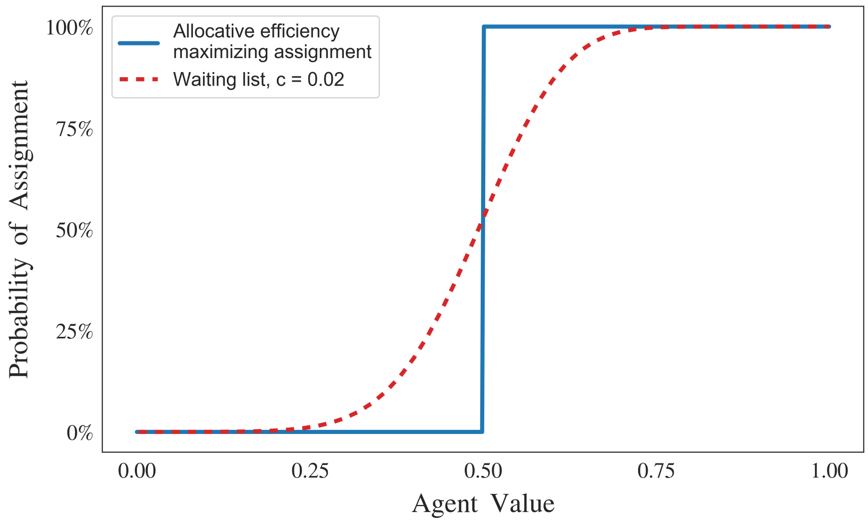

Figure 2: Assignment probabilities under the allocative efficiency-maximizing assignment and the waiting list assignment in Example 1.

tuate around market-clearing prices.[4] To quantify the impact of waiting time fluctuations on the economy, we evaluate the allocative efficiency of the waiting list assignment.[5]

A key quantity is the *adjustment size* Δ, which is defined to be the maximal change in an item's price (waiting cost) due to the addition of one agent to the item's queue. If waiting costs are linear, the adjustment size Δ equals the cost of waiting the expected arrival time of the least frequent item.[6]

Our main finding is that the allocative efficiency of the allocation generated by the waiting list is at least as high as the maximal possible allocative efficiency minus an efficiency loss that is roughly equal to the adjustment size Δ. Somewhat surprisingly, this bound holds for arbitrary agent value distributions. The bound is tight, and we construct a simple economy in which the allocative efficiency loss is approximately equal to the adjustment size Δ. We also identify economies in which a tighter bound holds and the allocative efficiency loss is small.

[4]Market-clearing prices maximize allocative efficiency, which may be distinct from social welfare. See the literature review below and Section 6 for a discussion of allocative efficiency and welfare.

[5]If all agent were to face market-clearing prices, the allocation would yield the maximal allocative efficiency. But, as in the example, fluctuations that result in waiting times that sufficiently diverge from market-clearing prices can lead to misallocation. Allocative efficiency loss quantifies the extent to which waiting time fluctuations meaningfully diverge from market-clearing prices.

[6]For example, suppose that public housing applicants incur a waiting cost of r USD per month because they need to pay higher open-market rent until receiving one of J different kinds of subsidized apartments. If each kind of apartment arrives once a month on average, then $\Delta = r$. If each kind of apartment arrives once a year on average, then $\Delta = 12 \cdot r$.

The adjustment size Δ decreases if all items arrive more frequently, or if agents incur a lower waiting cost per unit time.[7] Thus, our results imply that the waiting list asymptotically approaches the optimal allocative efficiency as item arrival rates increase or agent waiting costs decrease. In addition, we show that the waiting list asymptotically generates prices that are close to market-clearing prices with high probability. Together, these results imply that the outcome of the waiting list asymptotically approximates the competitive market outcome in terms of prices, assignment probabilities, and welfare.

As an illustration, consider the assignment of public housing apartments to applicants. The waiting list has compelling properties. It does not require any calibration, and is simple to implement. It allows the public housing authority to discover prices (waiting times) for each kind of apartment. Moreover, if demand changes (for example, some neighborhood becomes more attractive), prices in the waiting list automatically adjust. However, this price discovery also generates random fluctuations and misallocation. The allocative efficiency of the waiting list is roughly at least as high as optimally assigning all agents but making all agents wait for an additional arrival. The extent of potential misallocation depends on the environment: if each kind of apartment arrives frequently (say, weekly), then the adjustment size Δ is small and the waiting list generates close to the maximal allocative efficiency. But if each kind of apartment arrives rarely (say, annually), then the adjustment size Δ is large, prices can fluctuate considerably, and there can be significant allocative efficiency loss.

We prove our results and provide an intuition for why the efficiency loss is affected by the adjustment size through a novel connection between the waiting list and the stochastic gradient descent (SGD) algorithm. The SGD algorithm is an iterative optimization method that takes random steps that improve its objective in expectation.[8] When used for optimization, the SGD algorithm needs to gradually reduce the step size to zero as it approaches the optimal value (otherwise it can perpetually fluctuate around the optimal value without converging). We show that the run of the waiting list is equivalent to a run of the SGD algorithm that optimizes allocative efficiency, but with a step size that remains constant. The step size is exogenously given by the item arrival rates and the agent waiting costs, and is bounded by the adjustment size. Thus, prices under the waiting list tend to adjust toward market-clearing prices, but never converge due to the constant step size. Intuitively, the fluctuations and the efficiency loss are smaller if the adjustment size is smaller.

The connection to SGD allows us to employ new tools and avoid technical challenges in the analysis. A natural approach to analyzing the random price adjustment process is

[7]Market-clearing prices (denominated in units of utility) remain unchanged if all agents and items arrive proportionally more frequently, or if all agents incur a lower waiting cost per unit time.

[8]This algorithm is commonly used with great success and gained much attention in recent years because of its usefulness for training neural networks (see, e.g., LeCun et al., 2012).

to calculate the stationary distribution and the implied steady-state distribution of prices. However, when there are more than two item types the stationary distribution cannot be tractably calculated.[9] Instead, this paper takes a different technical approach, enabled by the connection to SGD. The proof leverages Lyapunov theory to obtain tractable bounds for allocative efficiency. By using a Lyapunov function, we are able to obtain a bound that does not require calculation of the stationary distribution.

Finally, we consider a price discovery process inspired by the waiting list and ask whether this simple price discovery process can perform well with a fixed adjustment size that is set appropriately. Consider a market with a finite time horizon in which items and agents arrive over time. A planner wishes to set monetary prices to maximize welfare, but has no knowledge of the distribution of agent preferences or the market-clearing prices. The waiting list suggests a simple pricing heuristic: increasing an item's price by a fixed adjustment size when it is demanded and decreasing an item's price by the adjustment size when it is supplied. The planner is not required to have any knowledge, except for calibrating the adjustment size according to the market's time horizon. We show that for an appropriately chosen adjustment size this simple heuristic results in an asymptotically optimal price discovery.

This simple heuristic exhibits a form of price rigidity. In markets with a finite time horizon, the optimal adjustment size needs to trade off between two different sources of loss. Roughly speaking, if prices are far from optimal, a large adjustment size helps the heuristic approach optimal prices faster. But when prices are close to optimal, a smaller adjustment size reduces price fluctuations by preventing the heuristic from reacting to noise. Knowledge of the market's time horizon allows the planner to choose an adjustment size that balances this trade-off. That is, the optimal adjustment size is flexible enough to allow the heuristic to learn and adapt, but rigid enough to prevent fluctuations.

Related Literature. This paper contributes to the growing literature on waiting lists as dynamic assignment mechanisms. Two important motivating applications are the assignment of public housing (Kaplan, 1984, 1988) and organ allocation (Zenios, 1999; Su and Zenios, 2004). The works of Bloch and Cantala (2017); Leshno (2017); Arnosti and Shi (2020); Baccara et al. (2018); Che and Tercieux (2021) develop tractable theoretical analysis of dynamic assignment problems. The previous theoretical literature restricts attention to markets with at most two kinds of items because the theoretical analysis relies on a calculation of the stationary distribution of the stochastic process, which is not possible when there are three or more kinds of items. This paper takes a different technical approach, allowing us to consider economies with any number of items.

[9] For this reason, several papers were limited to considering economies with at most two kinds of items (see, e.g., Baccara et al., 2018; Leshno, 2017; Bloch and Cantala, 2017).

A growing empirical literature complements the theory and studies the dynamic assignment of public housing apartments (Waldinger, 2021; van Dijk, 2019), transplant organs (Agarwal et al., 2021, 2020), cab rides (Buchholz et al., 2020), and hunting licences (Verdier and Reeling, 2022).

Several papers are concerned with the stochasticity in dynamic two-sided trading markets in order to optimize clearing timing decisions (Mendelson, 1982; Kelly and Yudovina, 2018; Loertscher et al., 2018). These papers also restrict attention to either a single asset or a binary type space.

We find that the waiting list's price discovery process generates waiting times that fluctuate around market-clearing prices. Although market-clearing prices maximize allocative efficiency, in various circumstances other rationing mechanisms may be preferable (Weitzman, 1977). Akbarpour et al. (2020); Dworczak et al. (2021) consider allocation mechanisms when agents may differ in their marginal utility of money and find that the optimal allocation can incorporate price controls. Hartline and Roughgarden (2008); Condorelli (2013) find that the allocation that maximizes consumer surplus may be a random lottery. These papers leverage the technical approach of Myerson (1981), and thus their analysis restricts attention to single-dimensional heterogeneity (e.g., agents that differ in their value for the quality of a good). In this paper we allow for any number of goods and heterogeneous preferences, and focus on price discovery in the waiting list.

Our paper connects to the large literature on strategic queueing, going back to Naor (1969). See Hassin and Haviv (2003); Hassin (2016) for a comprehensive review. This literature highlights strategic choices, externalities, and the inefficiencies caused by fluctuations in queue length. The present paper is distinct from this literature in that it considers the waiting list as an assignment mechanism and studies the matching between heterogeneous agents and heterogeneous items.

A growing literature studies matching in queues under nonstrategic stochastic arrivals. See Weiss (2021) for a review. A Markovian characterization of stochastic matching models is developed in Caldentey et al. (2009), Adan and Weiss (2012), and Adan et al. (2018), but Adan et al. (2018) conjectures that calculating the stationary distribution is computationally hard.

This paper also relates to papers that study the convergence of tâtonnement processes using gradient descent. Numerous paper analyze these processes in markets with multiple goods (Uzawa, 1960; Cole and Fleischer, 2008; Cheung et al., 2018, 2019) and in congestion or transportation settings (Powell and Sheffi, 1982; Correa and Stier-Moses, 2010). Bubeck et al. (2019) and Roth et al. (2020) leverage SGD-type algorithms to develop dynamic pricing algorithms that efficiently learn. These papers consider the problem of a planner that needs to learn prices and identify processes that converge to optimal prices. By contrast, under the waiting list prices never converge. A key distinction is that under the waiting list the adjustment size is fixed, while processes that converge to optimal prices must shrink the step size over time.

Organization of the Paper. Section 2 presents the model. The main results are stated in Section 3. Section 4 provides the intuition and proofs of our main results for linear waiting costs. The proof for general waiting costs is in Appendix A. Section 5 analyzes an SGD pricing heuristic that mimics the price adjustments of the waiting list and shows that such heuristics perform well in finite-horizon economies if the adjustment size is appropriately set. Section 6 concludes with a discussion of our modeling choices.

Appendix B contains all other omitted proofs. Appendix C provides additional results: calculating the adjustment size under nonlinear waiting costs, and providing explicit convergence rates.

2 Model

We study an infinite-horizon economy[10] in which agents and items arrive randomly over time. We describe the economy, set benchmarks for allocative efficiency, and describe the waiting list assignment.

Economy. We consider a market in which items and unit-demand agents arrive over time. Agents arrive according to a Poisson process with rate λ. Each agent has a type θ drawn independently according to a probability distribution F over the set of types Θ. We assume that Θ is a compact subset of a Euclidean space, and allow for both finitely many agent types as well as a continuum of agents.

Items arrive according to a Poisson process with total rate μ normalized to 1. The agent and item arrival processes are independent. Each arriving item is of a type $j \in \mathcal{J} = \{1, 2, \ldots, J\}$. An item is of type j with probability $\mu_j > 0$, where $\sum_{j\in\mathcal{J}} \mu_j = \mu = 1$. Denote $\mu_{\min} \triangleq \min_{j\in\mathcal{J}} \mu_j > 0$. We define an auxiliary item type $\emptyset$, which denotes being unassigned and use $\mathcal{J}_\emptyset \triangleq \mathcal{J} \cup \{\emptyset\}$.

The value an agent of type $\theta \in \Theta$ obtains when assigned an item of type $j \in \mathcal{J}_\emptyset$ is given by $v(\theta, j)$, where we normalize $v(\theta, \emptyset) = 0$. Agents' utilities are quasi-linear in waiting costs; an agent of type θ that is assigned an item of type j after waiting w units of time receives a utility of

$$u_\theta(j, w) \triangleq v(\theta, j) - c(w)\,,$$

[10] A finite-horizon economy is analyzed in Section 5.

where $c(w)$ is the cost of waiting w units of time.

A market is specified by $\left[\Theta, \mathcal{J}, F, \lambda, \{\mu_j\}_{j\in\mathcal{J}}, v(\cdot,\cdot), c(\cdot)\right]$. We say that there are finitely many agent types if F corresponds to finitely many atoms. A market with finitely many agent types can be specified by $\left[\Theta, \mathcal{J}, \boldsymbol{\lambda}, \boldsymbol{\mu}, \mathbf{v}, c(\cdot)\right]$ with $\boldsymbol{\lambda} \in \mathbb{R}_{++}^{|\Theta|}$, $\boldsymbol{\mu} \in \mathbb{R}_{++}^{|\mathcal{J}|}$, and $\mathbf{v} \in \mathbb{R}^{|\Theta|\times|\mathcal{J}|}$.

We make the following technical assumptions. We assume that for each $j \in \mathcal{J}$, $v(\theta, j)$ is continuous in θ and bounded from above by $v_{\max} \in \mathbb{R}$. We assume that the waiting cost function is smooth, strictly increasing, weakly concave,[11] and $c(0) = 0$, $\lim_{w\to\infty} c(w) = \infty$.

To simplify notation, we consider an equivalent discrete-time process,[12] indexed by t, that records the sequence of arrivals. For each arrival epoch t, the indicator ξ_t equals one if the t-th arrival is an agent, and equals zero if the t-th arrival is an item. If $\xi_t = 1$, let θ_t denote the type of the agent arriving at t. If $\xi_t = 0$, let $j_t \in \mathcal{J}$ denote the type of the item arriving at t. To simplify notation, we also let $\theta_t = \emptyset$ and $v(\theta_t, \cdot) = 0$ for $\xi_t = 0$.

Assignments and Allocative Efficiency. An *assignment* η assigns each arriving item to at most one agent. The allocative efficiency of a matching is defined as the average item's value for the agent to whom it has been assigned. Formally, given assignment η, for each epoch t such that $\xi_t = 1$, let $\eta_t \in \mathcal{J}_\emptyset$ be the kind of item assigned under η to the agent of type θ_t that arrived in epoch t. Let $A(T) = \{t \leq T \mid \xi_t = 1\}$ be the set of epochs up to epoch T in which agents arrived. Allocative efficiency under η is defined as

$$W(\eta) = \liminf_{T\to\infty} \frac{1}{|A(T)|} \sum_{t\in A(T)} v(\theta_t, \eta_t)\,. \tag{1}$$

We restrict attention to assignments that satisfy a *no-Ponzi condition*. Loosely speaking, this condition ensures that the assignment is approximately valid if the market terminates at some large finite time.[13] Formally, let $\mathcal{R}_T(\eta)$ denote the number of agents and items that arrive by time T and are waiting to be assigned at time[14] T. The assignment η satisfies the no-Ponzi condition if there exists a finite $M \in \mathbb{R}$ such that $\mathcal{R}_T(\eta) < M$ for all T.

[11]Our results also extend to convex $c(w)$ such that both $c'(w)$ and $c''(w)$ are subexponential; i.e., there exists α such that $c'(w), c''(w) \leq e^{\alpha w}$ for all $w \geq 0$.

[12]The equivalence is due to the Arrival Theorem of Poisson-driven processes (see, e.g., Wolff, 1982).

[13]For example, in the market of Example 1 there exists an assignment that assigns all agents items by assigning agents items that are further and further into the future (and the queue grows arbitrarily long). However, if the market terminates at any finite time, only approximately half the agents can be assigned items. The no-Ponzi condition rules out such problematic assignments.

[14]In other words, $\mathcal{R}_T(\eta)$ counts the number of agents who arrive before time T and are assigned under η an item that arrives after time T, plus the number of items that arrive before time T and are assigned to agents that arrive after time T.

Define optimal allocative efficiency to be

$$W^{\text{OPT}} = \mathbb{E}\left[\sup_{\eta \in H} W(\eta)\right] ,$$

where H is the set of no-Ponzi assignments and the expectation is taken over all possible realizations.

Assignment by a Waiting List. The waiting list holds a separate First-Come-First-Served queue for each item. An arriving agent observes the length of the queue for each item and chooses to join the end of one of the queues, or to leave immediately without being assigned an item (i.e., to balk). An agent who joins a queue will wait in that queue until receiving an item. When an item arrives, it is assigned to the agent at the head of its queue, if there is any; if the item's queue is empty, the item is discarded.

To formally describe the waiting list assignment, let $\mathbf{q} = (q_1, \ldots, q_J) \in \mathbb{Z}_+^J$ denote the state where there are q_j agents in the queue for item j. An arriving agent of type θ who observes $\mathbf{q}$ and chooses to join the queue for item j will wait a random amount of time w_j before receiving item $j \in \mathcal{J}_\emptyset$, and will receive an expected utility of $v(\theta, j) - \mathbb{E}[c(w_j)|q_j]$. Thus, the agent chooses to join the queue for item $a(\theta, \mathbf{q}) \in \mathcal{J}_\emptyset$ given by

$$a(\theta, \mathbf{q}) \in \operatorname*{argmax}_{j \in \mathcal{J}_\emptyset} \left\{ v(\theta, j) - p_j(q_j) \right\} , \tag{2}$$

where we define $p_j(q_j) \triangleq \mathbb{E}[c(w_j)|q_j]$. We refer to p_j as the *price* of item j. We allow agents to leave without joining any queue, and simplify notation by setting $p_\emptyset(\cdot) \equiv 0$. To further simplify notation, denote $\mathbf{p}(\mathbf{q}) \triangleq [p_1(q_1), \cdots, p_J(q_J)]$. Denote the queue lengths just before the t-th arrival by $\mathbf{q}_t$. That is, an agent that arrives at epoch t will face prices $\mathbf{p}_t = \mathbf{p}(\mathbf{q}_t)$, which depend on the current state of the queues $\mathbf{q}_t$. No queue length can ever exceed $q_{\max} \triangleq \max_{j \in \mathcal{J}} \ p_j^{-1}(v_{\max})$.

Given a realization, let η^{WL} denote the assignment induced by the waiting list. Under our assumptions, η^{WL} satisfies the no-Ponzi condition.[15] We denote the expected allocative efficiency of the waiting list by

$$W^{\text{WL}} \triangleq \mathbb{E}\left[W(\eta^{\text{WL}})\right] .$$

We refer to $W^{\text{OPT}} - W^{\text{WL}}$ as the *allocative efficiency loss*, or loss for short.

[15] Because no queue length can ever exceed $q_{\max} \triangleq \max_{j \in \mathcal{J}} \ p_j^{-1}(v_{\max})$ we have that $\mathcal{R}_t(\eta^{\text{WL}}) \ \leq \ |\mathcal{J}| \cdot q_{\max}$.

2.1 The static assignment problem and prices

It will be helpful to consider a natural corresponding static assignment problem in which all items and agents are simultaneously assigned. To transform the dynamic assignment problem into a static assignment problem, consider a planner that accumulates all agents and items that arrive up to some large time T. At time T, the planner assigns all agents and items to maximize allocative efficiency. In expectation, the number of items of each kind that arrive up to time T is proportional to each item's arrival rate (and likewise for agents).[16]

We denote the maximal value per agent generated by the static assignment by W^* and refer to it as the optimal static allocative efficiency. The static assignment problem and W^* are given by problem (3):

$$\begin{aligned} W^* = \max_{\{x_{\theta j}\}_{\theta\in\Theta, j\in\mathcal{J}}} \quad & \sum_{j\in\mathcal{J}} \int_\Theta x_{\theta j}\, v(\theta, j)\, dF(\theta) & & \\ \text{subject to} \quad & \sum_{j\in\mathcal{J}} x_{\theta j} \leq 1,\ x_{\theta j} \in [0,1] & \forall \theta \in \Theta & \qquad (3) \\ & \int_\Theta \lambda\, x_{\theta j}\, dF(\theta) \leq \mu_j & \forall j \in \mathcal{J}\ . & \end{aligned}$$

In problem (3), $x_{\theta j}$ is the share of agents of type θ that are assigned item j. The problem's objective is the average agent's value for their assigned item. The first constraint requires that the shares $x_{\theta j}$ be well defined. The second constraint is the resource constraint; it requires that the expected number of j items arriving per unit time be at least as large as the expected number of agents that arrive per unit time and that are assigned item j. The following proposition shows that optimal allocative efficiency is identical under the static and dynamic problems.

Proposition 1. *The optimal allocative efficiency for the dynamic assignment problem is*

$$W^{\text{OPT}} = W^*\ ,$$

where W^* is the optimal static allocative efficiency.

To gain intuition for Proposition 1, observe that $W^{\text{OPT}} \leq W^*$ because the static problem ignores the constraints imposed by the arrival order, and is therefore a relaxation of the dynamic problem.[17] To see that $W^{\text{OPT}} \geq W^*$, let $\mathbf{x}^*$ be an optimal solution to the static

[16] For any finite time T, the realized number of arrivals is stochastic, but the stochastic error becomes negligible as $T \to \infty$ by the central limit theorem.

[17] The no-Ponzi condition ensures that the long-run allocative efficiency is approximated by the allocative efficiency up to finite time T as $T \to \infty$.

assignment problem (3) and consider the dynamic assignment η that assigns an arriving agent θ the next unclaimed arrival of item j with probability[18] $x^*_{\theta j}$. The proof of Proposition 1 is in Appendix B.1.

The static assignment problem is also helpful for providing us with market-clearing prices. We use $\mathcal{P}^*$ to denote the set of market-clearing prices for the static assignment problem. Standard results (e.g., Bertsekas, 1981; Demange et al., 1986) imply the existence of market-clearing prices $\mathbf{p}^* \in \mathcal{P}^*$ that induce an optimal static assignment $\mathbf{x}^*$. Such market-clearing prices $\mathbf{p}^*$ can be used to maximize allocative efficiency in the dynamic problem: we can generate an optimal dynamic assignment by asking each arriving agent to join the queue for an item $j \in \mathcal{J}_\emptyset$ that maximizes $v(\theta, j) - p^*_j$.

3 Efficiency of the Price Discovery Process

Prices under the waiting list are not set by a planner or a market-clearing condition. Instead, prices change over time according to a tâtonnement-like process: the price (waiting cost) of item j increases when an agent chooses to join queue j, and the price of item j decreases when item j arrives and one agent is removed from queue j. As Example 1 illustrates, this price discovery process does not necessarily converge to market-clearing prices $\mathbf{p}^*$, or to any single price. Agents arriving to the waiting list may observe prices that differ from equilibrium prices, possibly resulting in lower allocative efficiency.

A natural approach to analyzing the waiting list is to calculate the distribution of prices agents face by calculating the stationary distribution of the waiting list. However, this stationary distribution is not tractable when $|\mathcal{J}| \geq 3$ (i.e., when there are more than 3 types of items).[19]

We therefore take a different approach that allows us to analyze general markets with any number of items, a general (possibly continuous) distribution of agent types, and non-linear waiting costs. Our analysis shows that the following attribute plays a central role in determining allocative efficiency:

Definition 1. *The* adjustment size Δ *is the maximal change in price due to a single arrival, and is given by*

$$\Delta \triangleq \max_{j \in \mathcal{J}} \max_{1 \leq q \leq q_{\max}} \big(p_j(q) - p_j(q-1)\big).$$

[18] If the resource constraint for item j holds with equality under $\mathbf{x}^*$, the number of agents waiting to be assigned item j evolves like the number of agents in a queue with arrival and departure rates equal to μ_j. Because this queue is critical, it violates the no-Ponzi condition. The proof shows that we can achieve approximately the same allocative efficiency with assignments that satisfy the no-Ponzi condition, for example by imposing a limit on the number of agents that can be waiting for item j.

[19] Because of this limitation, previous papers that relied on calculation of the stationary distribution were limited to models with at most two kinds of items (e.g., Leshno, 2017; Bloch and Cantala, 2017; Baccara et al., 2018).

In other words, each arrival of an item j reduces the price of item j by at most Δ. Each arrival of an agent who joins the queue of item j increases the price of item j by at most Δ. If waiting costs are linear $c(t) = c \cdot t$, we have $\Delta = c/\mu_{\min}$; that is, Δ is the expected cost of waiting for a single arrival of the least frequent item.

Our main result provides a lower bound for the allocative efficiency under the waiting list. Theorem 1 bounds $W^{\text{OPT}} - W^{\text{WL}}$, which we refer to as the allocative efficiency loss under the waiting list. This bound holds for any distribution of agent preferences.

Theorem 1. *The allocative efficiency under the waiting list is*

$$W^{\text{WL}} \geq W^{\text{OPT}} - \frac{\lambda + 2}{2\lambda}\Delta\,. \tag{4}$$

Theorem 1 states that allocative efficiency under the waiting list is at least the optimal allocative efficiency minus a loss that depends on the adjustment size. If waiting costs are linear and the arrival rate of agents is twice as high as the arrival rate of items ($\lambda = 2, \mu = 1$), then the allocative efficiency loss is bounded by the cost of waiting for a single arrival of the least frequent item. The loss will be small if all items arrive frequently and the cost of waiting for another arrival is small, but the loss can be large if items arrive infrequently.

As an illustration, consider the economy of Example 1. Recall that the distribution of agent values is $U[0, 1]$ and half the agents can be assigned an item, implying that the market-clearing price is $p^* = 0.5$. If $c(w) = 0.02w$, the market-clearing price is equivalent to the expected waiting time of an agent who joins a queue with 24 agents, as $\mathbb{E}[c(w)|24] = 0.5$. Each arrival or departure of an agent changes the price by $\Delta = 0.02$, which is 4% of p^*. By contrast, if the waiting cost is $\hat{c}(w) = 0.1w$, the market-clearing price is the expected waiting time of an agent who joins a queue with 4 agents, as $\mathbb{E}[\hat{c}(w)|4] = 0.5$. In this case, each arrival or departure of an agent changes the price by $\hat{\Delta} = 0.1$, which is 20% of p^*. Figure 3 presents the assignment probabilities under $c(\cdot)$ and $\hat{c}(\cdot)$, showing more misallocation under $\hat{c}(\cdot)$.

The proof of Theorem 1 in Section 4 shows that price adjustment in the waiting list is identical to the price adjustment of the stochastic gradient descent (SGD) optimization algorithm, which is an iterative optimization algorithm commonly used for training neural networks (see, e.g., LeCun et al., 2012). The SGD algorithm searches for an optimal solution by taking many iterative steps that each improve its objective in expectation. Each arrival moves prices in a random direction: a price increase if the arrival is an agent who chooses to join a queue, or a price decrease if the arrival is an item that removes an agent from a queue. But the expected change in an item's price is proportional to the difference between supply and demand for the item given current prices. In other words, in expectation each arrival moves prices toward market-clearing prices. This makes the price adjustment in the

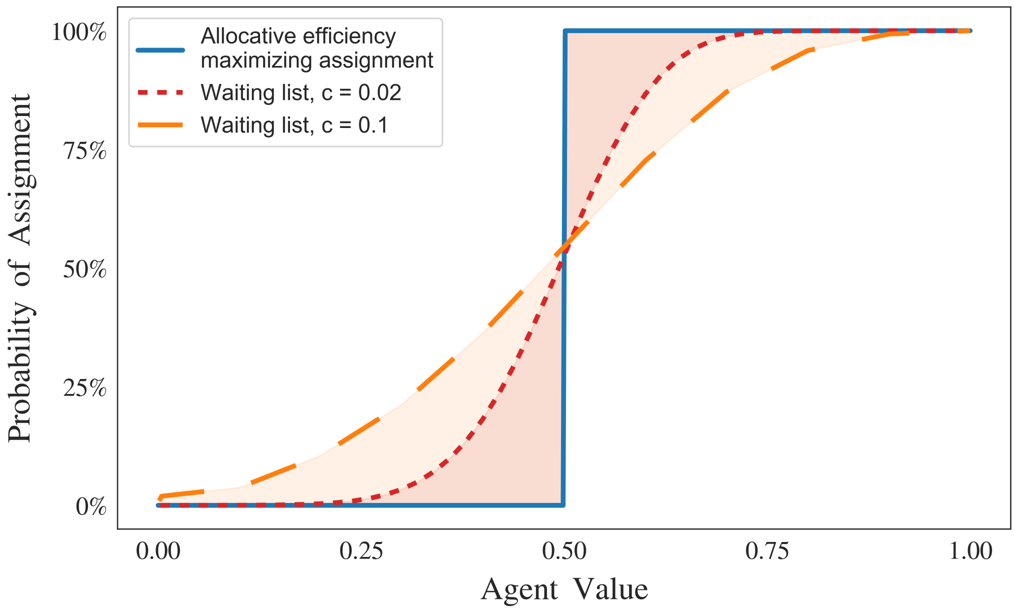

Figure 3: Assignment probabilities under $c(w) = 0.02w$ and under $\hat{c}(w) = 0.1w$.

waiting list equivalent to a run of the SGD algorithm.

The adjustment size corresponds to the step size of the SGD algorithm. When used for optimization, the SGD algorithm is run with a step size that decreases to zero over time. By contrast, the adjustment size in the waiting list is determined exogenously by the item arrival rates and agent waiting costs. Because the adjustment size is fixed and bounded away from zero, the SGD algorithm never converges.

Theorem 1 helps identify situations for which public housing assignment through the waiting list may perform poorly. The adjustment size is determined by the agents' cost of waiting, as well as the frequency of apartment arrivals. Thus, the allocative efficiency loss from the fluctuating waiting times will be small when apartments arrive frequently or when waiting costs per unit time are small. In such situations, the queues for each item will typically hold many agents, and relative waiting cost fluctuations will be small. The loss may be large if apartments arrive infrequently, or if waiting costs per unit time are significant.

The following theorem provides a tighter bound under a few additional assumptions. It shows that if all market-clearing prices are strictly positive, then the allocative efficiency loss is bounded by the adjustment size Δ plus an exponentially small loss.

Theorem 2. *Given an arbitrary distribution of agent types F and arbitrary item and agent arrival rates $\{\mu\}_{j\in\mathcal{J}}, \lambda$, consider market ℓ with linear waiting cost $c(w) = c_\ell \cdot w$. Let*

$W^{\mathrm{WL}}(c_\ell)$ denote the allocative efficiency of the waiting list for market ℓ, and let $\Delta_\ell = c_\ell/\mu_{\min}$ denote the adjustment size for market ℓ.

Suppose that $\lambda \geq 1$ and $p_j^* > 0$ for any item $j \in \mathcal{J}$ under any market-clearing prices $\mathbf{p}^*$. Then we have that

$$W^{\mathrm{WL}}(c_\ell) \geq W^{\mathrm{OPT}} - \Delta_\ell - \varepsilon$$

and there exist $\alpha, \beta, c_0 > 0$, such that for any $c_\ell < c_0$ we have $\varepsilon < \beta \cdot e^{-\alpha/\Delta_\ell}$.

The proof of Theorem 2 can be found in Appendix B.2. The remainder of this section provides additional results that further explore when the allocative efficiency loss is small or large.

3.1 Asymptotic optimality

The following corollaries of Theorem 1 show that the allocative efficiency under the waiting list approaches its optimal value as waiting costs per unit time become smaller or item arrivals become more frequent.

Corollary 1. *Given an arbitrary distribution of agent types F and arbitrary item and agent arrival rates $\{\mu\}_{j\in\mathcal{J}}, \lambda$, consider market ℓ with linear waiting cost $c(w) = c_\ell \cdot w$. Let $W^{\mathrm{WL}}(c_\ell)$ denote the allocative efficiency of the waiting list for market ℓ. We have that*

$$W^{\mathrm{WL}}(c_\ell) \xrightarrow[c_\ell \to 0]{} W^{\mathrm{OPT}}.$$

Observe that W^{OPT} and the market-clearing prices $\mathbf{p}^*$ are independent of c_ℓ. As $c_\ell \to 0$, agents become more patient and the adjustment size $\Delta = c_\ell/\mu_{\min}$ tends to zero. Thus, by Theorem 1 the loss tends to zero. Intuitively, the length of queue j that is required to generate the optimal price is $q_j^* = p_j^*/c_\ell$, which increases as $c_\ell \to 0$. When queue j is longer, the price fluctuations due to agents joining and leaving the queue become relatively small.

Corollary 2. *Given an arbitrary distribution of agent types F, arbitrary item and agent arrival rates $\{\mu\}_{j\in\mathcal{J}}, \lambda$, and waiting cost $c(\cdot)$, consider market n in which the agent arrival rate is $n \cdot \lambda$ and the arrival rate of item j is $n \cdot \mu_j$. Let $W^{\mathrm{WL}}(n)$ denote the allocative efficiency of the waiting list for market n. We have that*

$$W^{\mathrm{WL}}(n) \xrightarrow[n \to \infty]{} W^{\mathrm{OPT}}.$$

Intuitively, if the market thickens in the sense that arrivals of agents and items become more frequent, the expected cost of waiting for a single arrival becomes smaller. For example, if $c(w) = c \cdot w$ and $\mu_{\min} \to \infty$, then $\Delta = c/\mu_{\min} \to 0$.

3.2 Lower bound for the allocative efficiency loss

A natural question is whether the bound given in Theorem 1 is tight. This section gives a lower bound for the allocative efficiency loss by constructing an economy in which the loss is approximately Δ.

Proposition 2. *For any $J \in \mathbb{N}, \Delta > 0$ and $\varepsilon > 0$, there exists a market with J items in which*

$$W^{\mathrm{WL}} \leq W^{\mathrm{OPT}} - \Delta + \varepsilon \, .$$

The proof of Proposition 2 is in Appendix B.3. It shows the loss under the waiting list is large for the following market.

Example 2. *Consider a market in which $\Theta = \mathcal{J}$; that is, the set of items is $\mathcal{J} = \{1, 2, \ldots, J\}$ and there is a corresponding agent type for each item type. The distribution of agent types and the distribution of item types are uniform, i.e., $\mathbb{P}(\theta = j) = \mu_j = 1/J$, $\forall j \in \mathcal{J}$. The total agent arrival rate is $\lambda = 1$, and the waiting cost is linear, i.e., $c(w) = c \cdot w$. The value of agent θ for item j is*

$$v(\theta, j) = \begin{cases} \gamma & \text{if } \theta = j \, , \\ 0 & \text{if } \theta \neq j \, . \end{cases}$$

In the market of Example 2, any prices $\mathbf{p}$ where $0 \leq p_j \leq \gamma$ are market-clearing prices. When the price exceeds γ agents leave without being assigned, causing a loss equal to γ. The length of queue j follows a reflected unbiased random walk over $0, 1, 2, \ldots, \lceil \gamma/\Delta \rceil$, and an arriving agent is equally likely to observe either of these queue lengths. Thus, the probability that agents observe a price equal to $\Delta \cdot \lceil \gamma/\Delta \rceil > \gamma$ is roughly Δ/γ and the expected loss is roughly $\gamma \cdot \Delta/\gamma = \Delta$.

The example of Proposition 2 can be changed slightly to show that the bound in Theorem 2 is tight up to exponentially small terms. Although the market of Example 2 does not satisfy the condition that $p_j^* > 0$ under any market-clearing prices, we can create a similar market that satisfies this condition by slightly increasing the agent arrival rate by adding a small arrival rate of agents θ' that have a value $v(\theta', j) = \Delta$ for all goods $j \in \mathcal{J}$. The resulting market will have roughly the same allocative efficiency, and strictly positive prices for all items.

3.3 Small allocative efficiency loss with finitely many agents

The connection between the waiting list and the SGD optimization algorithm allows us to further explore the parameters of the economy that determine the magnitude of the allocative efficiency loss. In the market in Example 2, there is a wide range of market-clearing prices. Under the waiting list, prices fluctuate widely (see the proof of Proposition 2), resulting in frequent misallocation. However, in markets with finitely many agent types, multiplicity of market-clearing prices occurs only under a knife-edge set of parameters $(\boldsymbol{\lambda}, \boldsymbol{\mu}, \mathbf{v})$.

Lemma 1. *A market with finitely many agent types* $\left[\Theta, \mathcal{J}, \boldsymbol{\lambda}, \boldsymbol{\mu}, \mathbf{v}, c(\cdot)\right]$ *has unique market-clearing prices for an open and dense set of vectors* $(\boldsymbol{\lambda}, \boldsymbol{\mu}, \mathbf{v}) \in \mathbb{R}_{++}^{|\Theta|} \times \mathbb{R}_{++}^{|\mathcal{J}|} \times \mathbb{R}^{|\Theta|\cdot|\mathcal{J}|}$.

The following result shows that the waiting list performs well in a market with finitely many agent types that has unique market-clearing prices.

Theorem 3. *Given* $(\boldsymbol{\lambda}, \boldsymbol{\mu}, \mathbf{v}) \in \mathbb{R}_{++}^{|\Theta|} \times \mathbb{R}_{++}^{|\mathcal{J}|} \times \mathbb{R}^{|\Theta|\cdot|\mathcal{J}|}$*, consider market* ℓ *with linear waiting cost* $c(w) = c_\ell \cdot w$*. Assume that given* $(\boldsymbol{\lambda}, \boldsymbol{\mu}, \mathbf{v})$ *the market has unique market-clearing prices. Let* Δ_ℓ *denote the adjustment size for market* ℓ*. Let* $W^{\mathrm{WL}}(c_\ell)$ *denote the allocative efficiency of the waiting list for market* ℓ*. Then there exists* $\alpha, \beta, c_0 > 0$*, such that for any* $c_\ell < c_0$ *we have that*

$$W^{\mathrm{WL}}(c_\ell) \geq W^{\mathrm{OPT}} - \beta \cdot e^{-\alpha/\Delta_\ell} \, .$$

The intuition for Theorem 3 is as follows. If a market with finitely many agent types has unique market-clearing prices, the prices are robust in the following sense: as long as prices are within some $\delta > 0$ distance from the unique market-clearing prices $\mathbf{p}^*$, agents choose items they can be assigned under an efficient assignment. In other words, loss can occur only when agents observe a price that is more than δ away from $\mathbf{p}^*$. In addition, any price different from $\mathbf{p}^*$ causes significant imbalance between supply and demand in the static assignment problem. This implies that the price adjustment under the waiting list follows a biased random walk toward $\mathbf{p}^*$, and agents are unlikely to observe prices that are far from $\mathbf{p}^*$. Together, these facts imply that prices that cause misallocation are unlikely to occur.

3.4 Price concentration

The following result shows that the waiting costs borne by agents under the waiting list are approximately equal to market-clearing prices for the static assignment problem. We use $\mathrm{dist}(\mathbf{p}, \mathcal{P}^*) \triangleq \inf_{\mathbf{p}' \in \mathcal{P}^*} ||\mathbf{p} - \mathbf{p}'||_2$ to denote the minimal distance of a price to a market-clearing price, and use $B_\epsilon(\mathcal{P}^*) = \{\mathbf{p} \in \mathbb{R}_+^{|\mathcal{J}|} : \mathrm{dist}(\mathbf{p}, \mathcal{P}^*) < \epsilon\}$ to denote the set of prices that are within ϵ of some market-clearing price.

Proposition 3. *Given an arbitrary distribution of agent types* F *and arbitrary item and agent arrival rates* $\{\mu\}_{j \in \mathcal{J}}, \lambda$*, consider market* ℓ *with linear waiting cost* $c(w) = c_\ell \cdot w$ *and* *adjustment size Δ_ℓ. Let $\mathbf{p}(c_\ell)$ be a random variable generated by drawing the price observed by a randomly drawn agent arriving to the waiting list in market ℓ. Then for every $\epsilon > 0$ there exist $\alpha, \beta, c_0 > 0$, such that for any $c_\ell < c_0$ we have that*

$$\mathbb{P}\Big(\mathbf{p}(c_\ell) \notin B_\epsilon(\mathcal{P}^*)\Big) < \beta \cdot e^{-\alpha/\Delta_\ell}.$$

Proposition 3 and Corollary 1 together imply that if waiting costs per unit time are sufficiently small, the outcome of the waiting list approximates the competitive outcome of the static assignment problem. In particular, social welfare under the waiting list will be approximately equal to social welfare under the competitive outcome.

The proof of Proposition 3 is in Appendix B.5. Appendix C.2 gives explicit constants for the rate of convergence shown in Proposition 3.

An immediate implication of Proposition 3 is that items whose market-clearing price is strictly positive are rarely wasted.

Corollary 3. *Given a distribution of agent types F and item and agent arrival rates $\{\mu\}_{j\in\mathcal{J}}, \lambda$, consider market ℓ with linear waiting cost $c(w) = c_\ell \cdot w$ and adjustment size Δ_ℓ. Suppose that for item j we have that $p_j > 0$ under any market-clearing prices $\mathbf{p}$. Then there exist $\alpha, \beta, c_0 > 0$, such that, for every market with $c_\ell < c_0$, the probability that an arriving item j is discarded is at most $\beta \cdot e^{-\alpha/\Delta_\ell}$.*

4 Intuition and Proofs

Before providing the proof of Theorem 1, we show the connection between the waiting list and the stochastic gradient descent (SGD) optimization algorithm. This connection provides the intuition for Theorem 1 and inspired the technical approach of the proof. To formally show this connection, we use duality to formulate the static assignment problem as an optimization problem over possible prices. We then show that, in expectation, the run of the waiting list is a run of the SGD algorithm for this dual problem.

The Dual of the Static Assignment Problem. Consider the dual problem of the static assignment problem (3), which optimizes over possible prices. The following strong duality result is well known.[20]

Lemma 2. *[Monge–Kantorovich duality] The optimal value W^* of assignment problem (3)*

[20] Problem (3) is known as the (unbalanced) optimal transport problem, which has the strong duality property stated in Lemma 2. For further details, see, e.g., Galichon (2018).

coincides with the optimal value of the following dual optimization problem:

$$\underset{\mathbf{p}\geq\mathbf{0}}{\text{minimize}}\;\; h(\mathbf{p})\;,$$

where (here we let $p_\phi \triangleq 0$*)*

$$h(\mathbf{p}) \triangleq \int_\Theta \max_{j\in\mathcal{J}_\emptyset}\left[v(\theta,j)-p_j\right] dF(\theta) + \frac{1}{\lambda}\sum_{j\in\mathcal{J}} \mu_j p_j\;. \tag{5}$$

Throughout the proof, we denote by $\mathbf{p}^*$ an optimal value for the dual problem, and also refer to $\mathbf{p}^*$ as optimal prices or market-clearing prices.

Relation to Stochastic Gradient Descent. The SGD optimization algorithm is an iterative optimization algorithm. Each step of the SGD algorithm is random, but in expectation each step is a move in a subgradient direction of the algorithm's objective function. Under the waiting list, each arrival is a random draw of an agent or an item, but the expected arrival corresponds to supply and demand given current prices. In particular, the stochastic price adjustment from one arrival corresponds to a step in the SGD algorithm that seeks to minimize the dual objective $h(\mathbf{p})$. The next lemma formalizes this.

Lemma 3. *If the system is in state* $\mathbf{q}_t$ *such that* $q_{j,t}\geq 1$ *for all* j*, the expected change to the queue length from a single arrival* $\mathbb{E}[\mathbf{q}_t-\mathbf{q}_{t+1}]$ *equals* $\frac{\lambda}{1+\lambda}$ *times a subgradient of the dual objective* $h(\mathbf{p}_t)$ *at* $\mathbf{p}_t=\mathbf{p}(\mathbf{q}_t)$*.*

Indeed, observe that the expected adjustment to the length of queue j from a single arrival is

$$\begin{aligned}\mathbb{E}[q_{j,t+1}-q_{j,t}] &= \mathbb{E}\left[\mathbb{1}_{\{\xi_t=1,\,a(\theta_t,\mathbf{q}_t)=j\}} - \mathbb{1}_{\{\xi_t=0,\,j_t=j\}}\right] \\ &= \frac{\lambda}{1+\lambda}\int_\Theta \mathbb{1}_{\{j=\mathrm{argmax}_{j\in\mathcal{J}_\emptyset}\{v(\theta,j)-p_j(q_{j,t})\}\}} dF(\theta) - \frac{1}{1+\lambda}\mu_j\;,\end{aligned} \tag{6}$$

and that $\frac{1+\lambda}{\lambda}\mathbb{E}[q_{j,t+1}-q_{j,t}]$ is a subgradient of $h(\mathbf{p}_t)$ at $\mathbf{p}_t=\mathbf{p}(\mathbf{q}_t)$.

Therefore, on average, prices adjust in the right direction (but note that the actual adjustments are random and depend on the realization of the arrival). For a gradient descent optimization algorithm to converge, the step size must shrink when approaching the optimal value. Under the waiting list, the size of the adjustment is fixed and bounded by the step size Δ. Therefore, the price adjustment in the waiting list corresponds to the run of an SGD algorithm with a fixed step size that never converges. If the step size Δ is smaller, the algorithm will eventually fluctuate closer to the optimal value, leading to a smaller loss.

4.1 Proof of Theorem 1

The proof is based on a Lyapunov analysis. A Lyapunov function captures the "potential" in each state, which allows us to decompose the expected value from the next arrival into the objective, a change in potential, and a per-period loss.

We define the Lyapunov function to be the following quadratic function:

$$L(\mathbf{p}) = \frac{1}{2}\sum_{j\in\mathcal{J}} \mu_j p_j^2 \,.$$

The following lemma is the key step in the proof.

Lemma 4. *The following inequality holds:*

$$\mathbb{E}[v(\theta_t, a(\theta_t, \mathbf{q}_t))|\mathbf{q}_t] \geq \frac{\lambda}{\lambda+1}W^* - \underbrace{\frac{1}{\mu_{\min}\cdot\Delta}\big(L(\mathbf{p}_t) - \mathbb{E}[L(\mathbf{p}_{t+1})|\mathbf{q}_t]\big)}_{\text{(I) change in potential}} - \underbrace{\frac{2+\lambda}{2(1+\lambda)}\Delta}_{\text{(II) loss}} . \tag{7}$$

To interpret equation (7), observe that $\frac{\lambda}{\lambda+1}W^* = \frac{\lambda}{\lambda+1}W^{\text{OPT}}$ is the average per-arrival[21] value under the optimal assignment. Equation (7) shows that the waiting list achieves this value minus a change in potential and a per-period loss. Summing over many periods, the change in potential (I) forms a telescoping series, and therefore remains bounded. Therefore, as we average over many periods, we have that (I) tends to zero. The loss term (II) is uniformly bounded for any $\mathbf{p}_t$, which allows us to bound the loss without calculating the stationary distribution.

Proof of Theorem 1. Let $W_T(\eta^{\text{WL}})$ be the total value of items assigned to agents that arrive before epoch T. That is,

$$W_T(\eta^{\text{WL}}) = \sum_{t=1}^{T} v(\theta_t, a(\theta_t, \mathbf{q}_t)) \,,$$

where we abuse notation by defining $v(\theta_t, a(\theta_t, \mathbf{q}_t)) \equiv 0$ when[22] $\xi_t = 0$. Rescaling, we have that

$$W^{\text{WL}} = \frac{1+\lambda}{\lambda}\mathbb{E}\left[\liminf_{T\to\infty} \frac{W_T(\eta^{\text{WL}})}{T}\right].$$

Observe that

$$\mathbb{E}\left[W_T(\eta^{\text{WL}})\right] = \mathbb{E}\left[\sum_{t=1}^{T} v(\theta_t, a(\theta_t, \mathbf{q}_t))\right]$$

[21] Including both agent arrivals and item arrivals.
[22] That is, when an item arrives in epoch t.

$$
\begin{aligned}
&= \mathbb{E}\left[\mathbb{E}\left[\sum_{t=1}^{T} v(\theta_t, a(\theta_t, \mathbf{q}_t)) \mid \mathbf{q}_t\right]\right] \\
&= \mathbb{E}\left[\sum_{t=1}^{T} \mathbb{E}\left[v(\theta_t, a(\theta_t, \mathbf{q}_t)) \mid \mathbf{q}_t\right]\right] \\
&\geq \mathbb{E}\left[\sum_{t=1}^{T} \frac{\lambda}{1+\lambda} W^* - \frac{1}{\mu_{\min}\cdot\Delta}\left(L(\mathbf{p}_t) - \mathbb{E}[L(\mathbf{p}_{t+1} \mid \mathbf{q}_t])\right) - \frac{2+\lambda}{2(1+\lambda)}\Delta\right] \\
&= T\frac{\lambda}{1+\lambda} W^* - \frac{1}{\mu_{\min}\cdot\Delta}\left(L(\mathbf{p}_1) - \mathbb{E}[L(\mathbf{p}_{T+1})]\right) - T\frac{2+\lambda}{2(1+\lambda)}\Delta\,, \qquad (8)
\end{aligned}
$$

where the inequality follows from Lemma 4.

Lemma 8, stated in Appendix B.2, uses standard arguments to establish the following equality:

$$
W^{\mathrm{WL}} = \frac{1+\lambda}{\lambda} \lim_{T\to\infty} \frac{1}{T}\mathbb{E}\left[W_T(\eta^{\mathrm{WL}})\right]\,. \qquad (9)
$$

Equality (9) allows us to translate the bound for $\mathbb{E}[W_T(\eta^{\mathrm{WL}})]$ to a bound for W^{WL}. Plugging (8) into (9) and using Proposition 1, we obtain

$$
W^{\mathrm{WL}} \geq W^{\mathrm{OPT}} - \frac{2+\lambda}{2\lambda}\Delta,
$$

which completes the proof of Theorem 1. □

Proof of Lemma 4: We begin with a few definitions. Let $\mathbf{a}_t$ and $\mathbf{g}_t$ be the vectors of indicators representing the agents and items (goods) arriving at time t, respectively:

$$
\mathbf{a}_t \triangleq \mathbf{e}_{a(\theta_t,\mathbf{q}_t)}\xi_t\,, \qquad \mathbf{g}_t \triangleq \mathbf{e}_{j_t}(1-\xi_t)\,.
$$

That is, $a_{j,t} = \mathbb{1}_{\{\xi_t=1,\, a(\theta_t,\mathbf{q}_t)=j\}}$ and $g_{j,t} = \mathbb{1}_{\{\xi_t=0,\, j_t=j\}}$. Let the indicator $l_{j,t} \triangleq \max\{0, g_{j,t} - q_{j,t} - a_{j,t}\}$ denote whether an item of type j arrived at time t and was discarded.[23] Using this notation, we can write the dynamics governing the evolution of the length of queue j by

$$
q_{j,t+1} = [q_{j,t} + a_{j,t} - g_{j,t}]^+ = q_{j,t} + a_{j,t} - g_{j,t} + l_{j,t}\,, \quad \text{for each } j \in \mathcal{J}.
$$

[23] Recall that under our definition of the waiting list, an item is discarded if the item finds its corresponding queue to be empty when it arrives.

Step 1. We show that

$$\mathbb{E}\Big[v(\theta_t, a(\theta_t, \mathbf{q}_t)) - \langle \mathbf{p}_t, \mathbf{a}_t - \mathbf{g}_t\rangle \mid \mathbf{q}_t\Big] = \frac{\lambda}{1+\lambda} h(\mathbf{p}_t)\,. \tag{10}$$

Observe that

$$\begin{aligned}
&\mathbb{E}\Big[v(\theta_t, a(\theta_t, \mathbf{q}_t)) - \langle \mathbf{p}_t, \mathbf{a}_t - \mathbf{g}_t\rangle \mid \mathbf{q}_t\Big] \\
&= \mathbb{E}\Big[v(\theta_t, a(\theta_t, \mathbf{q}_t)) - \sum_{j\in\mathcal{J}} p_{j,t}(a_{j,t} - g_{j,t}) \mid \mathbf{q}_t\Big] \\
&= \mathbb{E}\Big[\max_{j\in\mathcal{J}_\emptyset}\ [v(\theta_t, j) - p_{j,t}] + \sum_{j\in\mathcal{J}} p_{j,t} g_{j,t} \mid \mathbf{q}_t\Big] \\
&= \mathbb{E}[\xi_t]\cdot\mathbb{E}\Big[\max_{j\in\mathcal{J}_\emptyset}\ [v(\theta_t, j) - p_{j,t}] \mid \xi_t = 1, \mathbf{q}_t\Big] + \mathbb{E}\Big[\sum_{j\in\mathcal{J}} p_{j,t} g_{j,t} \mid \mathbf{q}_t\Big] \\
&= \frac{\lambda}{1+\lambda}\int_\Theta \max_{j\in\mathcal{J}_\emptyset}\ [v(\theta_t, j) - p_{j,t}]\, dF(\theta) + \frac{1}{1+\lambda}\sum_{j\in\mathcal{J}} \mu_j p_{j,t} \\
&= \frac{\lambda}{1+\lambda} h(\mathbf{p}_t)\,.
\end{aligned}$$

Step 2. To simplify the exposition, we focus here on the case in which the waiting cost is linear. The proof for general waiting costs is in Appendix A.

Assume that the waiting cost takes a linear form, where $c(w) = c \cdot w$ for some $c > 0$. Thus, $c = \Delta \cdot \mu_{\min}$.

We show that:

$$\begin{aligned}
\mathbb{E}[v(\theta_t, a(\theta_t, \mathbf{q}_t))|\mathbf{q}_t] = \frac{\lambda}{1+\lambda} h(\mathbf{p}_t) - \frac{1}{c}\left(L(\mathbf{p}_t) - \mathbb{E}[L(\mathbf{p}_{t+1}) \mid \mathbf{q}_t]\right) \\
- \frac{1}{2}\sum_{j\in\mathcal{J}} \frac{c}{\mu_j}\mathbb{E}[(a_{j,t} - g_{j,t})^2 + l_{j,t}^2 \mid \mathbf{q}_t]
\end{aligned} \tag{11}$$

and

$$\frac{1}{2}\sum_{j\in\mathcal{J}} \frac{c}{\mu_j}\mathbb{E}[(a_{j,t} - g_{j,t})^2 + l_{j,t}^2 \mid \mathbf{q}_t] \le \frac{2+\lambda}{2(1+\lambda)}\Delta\,. \tag{12}$$

We first show equality (11). We have that the drift of Lyapunov function $L(\mathbf{p})$ in one period is

$$L(\mathbf{p}_t) - L(\mathbf{p}_{t+1})$$

$$
\begin{aligned}
&= \frac{1}{2}\sum_{j\in\mathcal{J}} \mu_j \left[p_{j,t}^2 - \left(p_{j,t} + \frac{c}{\mu_j}(a_{j,t} - g_{j,t} + l_{j,t}) \right)^2 \right] \\
&= \frac{1}{2}\sum_{j\in\mathcal{J}} \mu_j \left[p_{j,t}^2 - \left(p_{j,t} + \frac{c}{\mu_j}(a_{j,t} - g_{j,t}) \right)^2 - \frac{c^2}{\mu_j^2} l_{j,t}^2 - \frac{2c}{\mu_j} \left(p_{j,t} + \frac{c}{\mu_j}(a_{j,t} - g_{j,t}) \right) l_{j,t} \right] \\
&\overset{(a)}{=} \frac{1}{2}\sum_{j\in\mathcal{J}} \mu_j \left[p_{j,t}^2 - \left(p_{j,t} + \frac{c}{\mu_j}(a_{j,t} - g_{j,t}) \right)^2 - \frac{c^2}{\mu_j^2} l_{j,t}^2 \right] ,
\end{aligned}
$$

where equality (a) follows from the fact that $\left(p_{j,t} + \frac{c}{\mu_j}(a_{j,t} - g_{j,t}) \right) l_{j,t} \equiv 0$ for all $j \in \mathcal{J}$.

We further simplify the Lyapunov drift as follows:

$$
\begin{aligned}
-\frac{1}{c}\big(L(\mathbf{p}_t) - L(\mathbf{p}_{t+1})\big) &= -\frac{1}{2c}\sum_{j\in\mathcal{J}} \mu_j \left[p_{j,t}^2 - \left(p_{j,t} + \frac{c}{\mu_j}(a_{j,t} - g_{j,t}) \right)^2 \right] + \frac{1}{2}\sum_{j\in\mathcal{J}} \frac{c}{\mu_j} l_{j,t}^2 \\
&= \frac{1}{2}\sum_{j\in\mathcal{J}} \mu_j \left[\frac{2}{\mu_j}(a_{j,t} - g_{j,t}) p_{j,t} + \frac{c}{\mu_j^2}(a_{j,t} - g_{j,t})^2 \right] + \frac{1}{2}\sum_{j\in\mathcal{J}} \frac{c}{\mu_j} l_{j,t}^2 \\
&= \sum_{j\in\mathcal{J}} (a_{j,t} - g_{j,t}) p_{j,t} + \frac{1}{2}\sum_{j\in\mathcal{J}} \frac{c}{\mu_j} \left((a_{j,t} - g_{j,t})^2 + l_{j,t}^2 \right) \\
&= \langle \mathbf{p}_t, \mathbf{a}_t - \mathbf{g}_t \rangle + \sum_{j\in\mathcal{J}} \frac{c}{2\mu_j} \left((a_{j,t} - g_{j,t})^2 + l_{j,t}^2 \right) . \qquad (13)
\end{aligned}
$$

By taking the expectation conditional on $\mathbf{q}_t$ on both sides of equation (13), adding equation (10), and rearranging, we establish equality (11).

We proceed to show inequality (12). Observe that

$$
\frac{c}{2\mu_j} \left((a_{j,t} - g_{j,t})^2 + l_{j,t}^2 \right) = \begin{cases} \frac{c}{\mu_j} & \text{if } g_{j,t} = 1 \text{ and } q_j = 0 \,, \\ \frac{c}{2\mu_j} & \text{otherwise.} \end{cases}
$$

That is, the above term is equal to $c/\mu_j \leq \Delta$ for arrivals that correspond to an item that is discarded because its queue is empty, and equal to $c/2\mu_j \leq \Delta/2$ for all other arrivals. Note that the probability that an item is discarded is at most $1/(1+\lambda)$ (the probability that an item arrives). Therefore

$$
\begin{aligned}
&\frac{1}{2}\sum_{j\in\mathcal{J}} \frac{c}{\mu_j} \mathbb{E}[(a_{j,t} - g_{j,t})^2 + l_{j,t}^2 \mid \mathbf{q}_t] \\
&\leq \frac{\Delta}{2} + \frac{1}{1+\lambda}\frac{\Delta}{2} \\
&= \frac{2+\lambda}{2(1+\lambda)}\Delta \,.
\end{aligned}
$$

The lemma follows from equations (11) and (12) and the inequality $h(\mathbf{p}_t) \geq W^*$, which is implied by Lemma 2.

□

4.2 Proof of Theorem 3

Proof of Theorem 3. Consider a market with finitely many agent types and linear waiting costs $\left[\Theta, \mathcal{J}, \boldsymbol{\lambda}, \boldsymbol{\mu}, \mathbf{v}, c(\cdot)\right]$ that has unique dual price $\mathbf{p}^*$. To simplify notation, we use $\mathbf{p}_t$ to denote the random variable generated by drawing the price observed by a randomly drawn agent arriving to the waiting list. By the arrival theorem, the distribution of $\mathbf{p}_t$ is equal to the steady state distribution.

Step 1. We first establish that there exists a $\delta > 0$ such that the allocative efficiency loss is small if prices are within δ of $\mathbf{p}^*$ with high probability.

For each agent $\theta \in \Theta$, let $\mathcal{J}_\theta^* \triangleq \operatorname{argmax}_{j \in \mathcal{J}_\emptyset}\{v(\theta, j) - p_j^*\}$ be the set of items agent θ can optimally demand under prices $\mathbf{p}^*$. The following definition[24] of δ_1 ensures that for any $\mathbf{p}$ such that $\|\mathbf{p} - \mathbf{p}^*\|_\infty < \delta_1$ we have that $a(\theta, \mathbf{p}) \in \mathcal{J}_\theta^*$ for all $\theta \in \Theta$:

$$\delta_1 \triangleq \min_{\theta \in \Theta}\left\{\max_{j \in \mathcal{J}_\emptyset}\{v(\theta, j) - p_j^*\} - \max_{j \in \mathcal{J}_\emptyset \setminus \mathcal{J}_\theta^*}\{v(\theta, j) - p_j^*\}\right\} .$$

Let $\delta_2 = \min_{j \in \mathcal{J}}\{p_j^* \mid p_j^* > 0\}$, and set $\delta = \frac{1}{2}\min\{\delta_1, \delta_2\}$.

Step 2. Let $\kappa = \mathbb{P}\left(\|\mathbf{p}_t - \mathbf{p}^*\|_\infty > \delta\right)$. By Proposition 3, there exist α, β, c_0 such that $\kappa < \beta e^{-\alpha/\Delta_\ell}$ for any $c < c_0$. We choose $c_0 > 0$ to be sufficiently small so that $\beta e^{-\alpha/\Delta_\ell} < 1/2$. In addition, we choose $c_0 > 0$ to be sufficiently small so that $c_0/\mu_{\min} < \delta/2$, which implies that if at time t we have that $\|\mathbf{p}_t - \mathbf{p}^*\|_\infty < \delta$, then for any j such that $p_j^* > 0$ the j-th queue is not empty at epoch t (that is, $q_{jt} > 0$).

Step 3. Let $\mathbf{x}$ denote the assignment probabilities of agents arriving to the waiting list in states where $\|\mathbf{p}_t - \mathbf{p}^*\|_\infty < \delta$. By the arrival theorem, $\mathbf{x}$ is given by

$$\mathbf{x}_{\theta j} = \mathbb{E}\left[\mathbb{1}_{\{a(\theta, \mathbf{p}_t) = j\}} \mid \|\mathbf{p}_t - \mathbf{p}^*\|_\infty < \delta\right] .$$

Let $\mathbf{y}$ denote the assignment probabilities of the remaining agents, i.e.,

$$y_{\theta j} = \mathbb{E}\left[\mathbb{1}_{\{a(\theta, \mathbf{p}_t) = j\}} \mid \|\mathbf{p}_t - \mathbf{p}^*\|_\infty \geq \delta\right] ,$$

and denote $x_j = \sum_{\theta \in \Theta} \lambda_\theta x_{\theta j},\ y_j = \sum_{\theta \in \Theta} \lambda_\theta y_{\theta j}$.

[24] We follow the convention that if $\mathcal{J}_\emptyset \setminus \mathcal{J}_\theta^* = \emptyset$ then $\max_{j \in \mathcal{J}_\emptyset \setminus \mathcal{J}_\theta^*}\{v(\theta, j) - p_j^*\} = -\infty$.

Because all agents are eventually assigned, we have that $(1-\kappa)x_j + \kappa y_j \le \mu_j$ for every $j \in \mathcal{J}$. If an item j with $p^*_j > 0$ arrives in epoch t such that $\|\mathbf{p}_t - \mathbf{p}^*\|_\infty < \delta$, then the item is assigned to a waiting agent. Therefore, at most $\kappa \cdot \mu_j$ of j items are wasted, and for every $j \in \mathcal{J}$ such that $p^*_j > 0$ we have that $(1-\kappa)x_j + \kappa y_j \ge (1-\kappa)\mu_j$ or

$$x_j \ge \mu_j - \frac{\kappa}{1-\kappa} y_j \ge \mu_j - \frac{\kappa}{1-\kappa}\lambda \,. \tag{14}$$

Observe that for an optimal assignment $\mathbf{x}^*$ we can write

$$\begin{aligned}
W^* &= \sum_{\theta,j} \lambda_\theta x^*_{\theta j} v(\theta, j) \\
&= \sum_{\theta,j} \lambda_\theta x^*_{\theta j} (v(\theta,j) - p^*_j) + \sum_{\theta,j} \lambda_\theta x^*_{\theta j} p^*_j \\
&= \sum_{\theta,j} \lambda_\theta x^*_{\theta j} (v(\theta,j) - p^*_j) + \sum_{j} \left(\mu_j p^*_j + (x^*_j - \mu_j) p^*_j\right) \\
&= \sum_{\theta,j} \lambda_\theta x^*_{\theta j} (v(\theta,j) - p^*_j) + \sum_{j} \mu_j p^*_j \,,
\end{aligned}$$

where the last equality follows from complementary slackness. By the definition of δ, we have that for every θ, j such that $x_{\theta j} > 0$ it is the case that $j \in \mathcal{J}^*_\theta$, implying that

$$\sum_{\theta,j} \lambda_\theta x^*_{\theta j} (v(\theta,j) - p^*_j) = \sum_{\theta,j} \lambda_\theta x_{\theta j} (v(\theta,j) - p^*_j) \,.$$

Thus, we have

$$\sum_{\theta,j} \lambda_\theta x_{\theta j} (v(\theta,j) - p^*_j) = W^* - \sum_j \mu_j p^*_j \,. \tag{15}$$

Using (14),(15) we therefore have

$$\begin{aligned}
\mathbb{E}\Big[v(\theta_t, a(\theta_t, \mathbf{p}_t))\Big] &= (1-\kappa)\mathbb{E}\Big[v(\theta_t, a(\theta_t, \mathbf{p}_t)) \mid \|\mathbf{p}_t - \mathbf{p}^*\|_\infty < \delta\Big] \\
&\quad + \kappa \mathbb{E}\Big[v(\theta_t, a(\theta_t, \mathbf{p}_t)) \mid \|\mathbf{p}_t - \mathbf{p}^*\|_\infty \ge \delta\Big] \\
&\ge (1-\kappa) \sum_{\theta,j} \lambda_\theta x_{\theta j} v(\theta, j) \\
&= (1-\kappa) \left(\sum_{\theta,j} \lambda_\theta x_{\theta j} \big(v(\theta,j) - p^*_j\big) + \sum_j x_j p^*_j \right) \\
&= (1-\kappa) \left(W^* + \sum_j (x_j - \mu_j) p^*_j \right)
\end{aligned}$$

$$
\begin{aligned}
&\geq (1-\kappa)\left(W^* + \sum_j (\mu_j - \frac{\kappa}{1-\kappa}\lambda - \mu_j) p_j^*\right) \\
&\geq (1-\kappa)\left(W^* - J\frac{\kappa}{1-\kappa}\lambda\, v_{\max}\right) \\
&\geq W^* - \kappa W^* - \kappa J \lambda\, v_{\max}\,.
\end{aligned}
$$

We chose c_0 such that $\kappa < \beta e^{-\alpha/\Delta_\ell}$, and therefore we have that

$$
\begin{aligned}
W^{\mathrm{WL}}(c_\ell) &= \mathbb{E}\Big[v(\theta_t, a(\theta_t, \mathbf{p}_t))\Big] \\
&\geq W^* - \kappa\,(W^* + J\lambda\, v_{\max}) \\
&\geq W^* - \beta\,(W^* + J\lambda\, v_{\max})\, e^{-\alpha/\Delta_\ell}\ ,
\end{aligned}
$$

which completes the proof. □

5 SGD Pricing Heuristic with Optimal Adjustment Size

In this section we consider a planner who can set prices, but does not know the distribution of agent preferences or the market-clearing prices. The waiting list suggests a simple pricing heuristic: prices adjust over time, increasing an item's price when it is demanded and decreasing an item's price when it is supplied. We show that this simple pricing heuristic approximately maximizes allocative efficiency. This heuristic can be used by a planner that has minimal knowledge of the market: the planner only needs to set the adjustment size according to the market's time horizon.

Finite-Horizon Economy. The economy is as defined as in Section 2, except that the world ends after T arrival epochs and agents have no cost of waiting (i.e., $c(w) \equiv 0$). A feasible assignment η assigns each agent that arrives before time T to an item that arrives before time T (or leaves the agent unassigned). For notational convenience, we define the allocative efficiency of allocation η as the total value of assigned items, i.e.,

$$
W_T(\eta) = \sum_{t=1}^{T} v(\theta_t, \eta_t)\,. \tag{16}
$$

Define the optimal allocative efficiency to be

$$
W_T^{\mathrm{OPT}} = \mathbb{E}\left[\sup_\eta W_T(\eta)\right]\ ,
$$

where we take the supremum over all possible assignments and the expectation is taken over all possible realizations. The following proposition relates the optimal allocative efficiency in the finite-horizon market to the optimal allocative efficiency of the static assignment problem.

Proposition 4. *Let $W_T^* \triangleq \frac{\lambda}{\lambda+1} T \cdot W^*$, where W^* is the optimal static allocative efficiency. The optimal allocative efficiency in the finite-horizon economy is bounded by*

$$W_T^{\mathrm{OPT}} \leq W_T^* .$$

The proof of Proposition 4 is almost identical to the proof of Proposition 1 and is therefore omitted.

5.1 The SGD pricing algorithm.

We now formally define the SGD pricing algorithm, which is inspired by the waiting list. Instead of using waiting times, the SGD pricing algorithm uses monetary prices. Prices adjust with each arrival, replicating the price adjustment under the waiting list except that the adjustment size can be specified by the planner (in contrast to the exogenously given adjustment size of the waiting list).

With a slight abuse of notation, we use $\mathbf{p}_t$ to denote the monetary prices used by the planner in arrival epoch t. In an epoch t in which an agent arrives, (i.e., $\xi_t = 1$), the agent can choose any item $j \in \mathcal{J}_\emptyset$ at price[25] $p_{j,t}$. An agent of type θ_t chooses item type

$$a(\theta_t, \mathbf{p}_t) \in \operatorname*{argmax}_{j \in \mathcal{J}_\emptyset} \left\{ v(\theta, j) - p_{j,t} \right\} .$$

If $a(\theta_t, \mathbf{p}_t) = j \in \mathcal{J}$ the agent joins the queue to wait for item j (recall that waiting is costless). In an epoch t in which an item arrives (i.e., $\xi_t = 0$), denote the item's type by $j_t \in \mathcal{J}$. The item is assigned to some agent waiting in the queue for that type of item.[26] If there are no agents waiting for item j_t, the item is discarded.

Prices adjust with each arrival, as follows. Initialize the prices of all items to zero, i.e., $p_{1,j} = 0$, $\forall j \in \mathcal{J}$. Prices for epoch $t+1$ are adjusted according to the arrival in epoch t as follows:

$$p_{j,t+1} = \left[p_{j,t} + \Delta \cdot \mathbb{1}_{\{\xi_t = 1, j = a(\theta_t, \mathbf{p}_t)\}} - \Delta \cdot \mathbb{1}_{\{\xi_t = 0, j = j_t\}} \right]^+ . \tag{17}$$

That is, the price adjustment mimics the adjustment of waiting times under the waiting

[25] The price of being unassigned is $p_{\emptyset,t} = 0$ for all $t \leq T$.

[26] For concreteness, items are assigned to agents in the queue according to a First-Come First-Served order. Because waiting is costless, the choice of agent in the queue can be arbitrary.

list. If an agent chooses item j, the price of item j increases by a Δ increment. If item j arrives, the price of item j decreases by Δ (or stays at 0).

Let W_T^{WL} denote the allocative efficiency of the SGD pricing algorithm. Theorem 4 shows that this simple pricing algorithm performs well if the adjustment size Δ is set appropriately.

Theorem 4. *The allocative efficiency of the SGD algorithm with step size $\Delta = 1/\sqrt{T}$ satisfies*

$$W_T^{\mathrm{WL}}/T \geq W_T^*/T - \frac{1}{\sqrt{T}}\left(\frac{2+\lambda}{2+2\lambda} + J \cdot v_{\max}^2\right).$$

Setting the adjustment size $\Delta = 1/\sqrt{T}$ allows the pricing algorithm to obtain total allocative efficiency that is within $O(\sqrt{T})$ of W_T^*. Per-arrival allocative efficiency is within $O(1/\sqrt{T})$ of W_T^*/T. Devanur et al. (2019) show that the total allocative efficiency under any online pricing algorithm is at most $W_T^* - \Omega(\sqrt{T})$, implying that the SGD pricing algorithm is asymptotically optimal.[27]

We note two immediate extensions of Theorem 4. First, regardless of the choice of initial prices $\mathbf{p}_1 \in [0, v_{\max}]^J$, the allocative efficiency of the SGD pricing algorithm is within $O(\sqrt{T})$ of the optimal W_T^*. Second, setting the adjustment size to $\Delta = v_{\max}/\sqrt{T}$ can improve the allocative efficiency bound. We state both of these observation in the following corollary.

Corollary 4. *The allocative efficiency of the SGD algorithm with arbitrary initial prices $\mathbf{p}_1 \in [0, v_{\max}]^J$ and adjustment size $\Delta = v_{\max}/\sqrt{T}$ satisfies*

$$W_T^{\mathrm{WL}}/T \geq W_T^*/T - \frac{v_{\max}}{\sqrt{T}}\left((1+\mu_{\min}^{-1})J + \frac{2+\lambda}{2+2\lambda}\right).$$

Intuition and Proof. Before giving the proof of Theorem 4, we give a stylized explanation of the result. Roughly speaking, we can divide the run of the SGD pricing algorithm into two phases. First, the algorithm gradually raises prices to "learn" the market-clearing prices. Second, after some time the distribution of prices is close to the steady-state distribution of prices. In the second phase the algorithm incurs losses from price fluctuations (as discussed in Section 3). In the first phase prices are too low, allowing agents with low valuations to join the queue. This is wasteful because there is a limited supply of items that arrive by time T: although we always allow agents to join the queue for any item j, if too many agents choose item j, then at time T there will be agents waiting in queue j, and these agents will not be assigned any item.

[27] Given an arbitrary function $f(\cdot)$, and a positive function $h(\cdot)$, we say that $f(\rho) = O(h(\rho))$ if there exist positive constants α, ρ_0 such that $0 \leq f(\rho) \leq \alpha h(\rho)$ for any $\rho > \rho_0$. Similarly, we say that $f(\rho) = g(\rho) - \Omega(h(\rho))$ if there exist positive constants α, ρ_0 such that $0 \leq \alpha h(\rho) \leq g(\rho) - f(\rho)$ for any $\rho > \rho_0$.

A lower adjustment size Δ reduces the loss from fluctuations in the second phase. A higher adjustment size Δ helps the algorithm approach market-clearing prices faster and reduces the loss from the first phase. Setting the adjustment size to an intermediate value $\Delta = 1/\sqrt{T}$ allows the algorithm to balance the two sources of loss.

Proof. Let $\tilde{W}_T(\eta) \triangleq \sum_{t=1}^{T} v(\theta_t, \eta_t)$ denote the allocative efficiency of the SGD algorithm if all agents who choose an item are assigned; that is, $\tilde{W}_T(\eta)$ is the allocative efficiency of the SGD algorithm if there were no supply constraints. We separately bound $W_T^* - \tilde{W}_T^{\mathrm{WL}}$ and $W_T^{\mathrm{WL}} - \tilde{W}_T^{\mathrm{WL}}$.

From the analysis of Theorem 1, equation (8), we have:

$$\mathbb{E}\left[\tilde{W}_T(\eta^{\mathrm{WL}})\right] \geq T\frac{\lambda}{1+\lambda}W^* - \frac{1}{\Delta\mu_{\min}}\left(L(\mathbf{p}_1) - \mathbb{E}[L(\mathbf{p}_{T+1})]\right) - T\frac{2+\lambda}{2(1+\lambda)}\Delta\,.$$

Therefore we have:

$$W_T^* - \tilde{W}_T^{\mathrm{WL}} \leq \frac{1}{\Delta\mu_{\min}}\big(L(\mathbf{p}_1) - \mathbb{E}[L(\mathbf{p}_{T+1})]\big) + T\frac{2+\lambda}{2(1+\lambda)}\Delta\,.$$

Setting $\Delta = 1/\sqrt{T}$, and observing that $L(\cdot) \geq 0$ and $L(\mathbf{p}_1) = L(0) = 0$, we have that

$$W_T^* - \tilde{W}_T^{\mathrm{WL}} \leq \sqrt{T}\frac{2+\lambda}{2(1+\lambda)}\,. \tag{18}$$

To bound $\tilde{W}_T^{\mathrm{WL}} - W_T^{\mathrm{WL}}$, observe that this difference is the total value of items assigned to agents that are waiting in the queue in the final epoch T (when the economy ends). By definition, the number of agents waiting for item j at time T is $p_{j,T}/\Delta$. Since no agent will join queue j when $p_{j,t} > v_{\max}$, we have that $p_{j,t} \leq v_{\max}$ for all $j \in \mathcal{J}$ and $t \leq T$. Therefore,

$$\tilde{W}_T^{\mathrm{WL}} - W_T^{\mathrm{WL}} \leq \sum_{j\in\mathcal{J}} v_{\max}\cdot\frac{p_{j,T}}{\Delta} \leq J\cdot v_{\max}^2\cdot\sqrt{T}\,. \tag{19}$$

Combining (18) and (19) completes the proof. □

Proof of Corollary 4. The corollary follows from the proof of Theorem 4 by plugging $\Delta = v_{\max}/\sqrt{T}$ to obtain the following inequalities:

$$\begin{aligned} W_T^* - \tilde{W}_T^{\mathrm{WL}} &\leq \frac{1}{\Delta\mu_{\min}}\big(L(\mathbf{p}_1) - \mathbb{E}[L(\mathbf{p}_{T+1})]\big) + T\frac{2+\lambda}{2(1+\lambda)}\Delta \\ &\leq \frac{\sqrt{T}}{\mu_{\min}v_{\max}}\sum_{j\in\mathcal{J}}\mu_j p_{j,1}^2 + T\frac{2+\lambda}{2(1+\lambda)}\frac{v_{\max}}{\sqrt{T}} \end{aligned}$$

$$\leq \mu_{\min}^{-1} J v_{\max}\sqrt{T} + \frac{2+\lambda}{2+2\lambda} v_{\max}\sqrt{T}\,.$$

And,

$$\tilde{W}_T^{\mathrm{WL}} - W_T^{\mathrm{WL}} \leq \sum_{j\in\mathcal{J}} v_{\max}\cdot\frac{p_{j,T}}{\Delta} = \sum_{j\in\mathcal{J}} v_{\max}\cdot\frac{p_{j,T}}{v_{\max}/\sqrt{T}} \leq J\cdot v_{\max}\cdot\sqrt{T}\,.$$

Combining these inequalities and rearranging terms yields the desired result.

□

6 Conclusion

This paper analyzed the endogenous determination of waiting times under the waiting list. Waiting times continuously adjust according to a tâtonnement-like process, which is equivalent to running the SGD optimization algorithm with a fixed and exogenously given adjustment size. The resulting allocative efficiency is at least the maximal allocative efficiency minus a loss equal to the adjustment size, given by the marginal increase in expected waiting cost from having one more agent in the queue. Waiting lists approximately replicate the static competitive equilibrium outcome when the adjustment size is small, but can incur an additional loss from fluctuations when the adjustment size is large.

In practice, the adjustment size will be large in applications where the waiting cost per time unit is large or items arrive infrequently. In the context of public housing assignment, our results raise concerns about losses due to misallocation when applicants have urgent housing needs or when apartments arrive infrequently. We emphasize that such losses can occur even if there are many agents who seek housing. On the other hand, if housing projects are large and similar apartments arrive frequently, losses due to misallocation are small.

To focus our analysis on the endogenous determination of waiting times, we focused on the standard waiting list, which holds a separate queue for each item. This mechanism is commonly used, and is particularly suitable for our analysis because agents face a single choice that is analogous to the choice agents make in competitive equilibrium models. Insights from the analysis apply to other mechanisms as well. For example, changes to the queueing policy (e.g., Leshno, 2017) can effectively reduce the adjustment size and thus reduce misallocation. Similarly, giving agents infrequent updates can be seen as another method for reducing the effective adjustment size.

Our analysis abstracts away from potential distinctions between waiting times and monetary prices. If payments for items are transfers, competitive equilibria is welfare-maximizing (in addition to maximizing allocative efficiency). In applications such as public housing waiting lists, applicants bear a cost while they wait for an apartment, e.g., higher rent payments. Potential applicants bear this waiting cost whether or not they participate in the waiting list. Each assigned apartment reduces the waiting cost of one agent. Thus, total waiting costs are constant under any assignment that assigns all available apartments, and waiting times are transfers that do not affect social welfare. Our analysis shows that if all items have a positive price, the waiting list assigns almost all available items. However, we note that allocative efficiency maximization is distinct from welfare maximization in applications where agents have an outside option that reduces their waiting costs (for example, visitors to a theme park can choose to watch a show instead of queueing for a ride).

Finally, we consider a simple pricing heuristic inspired by the waiting list. We show that a simple SGD pricing heuristic with an appropriate adjustment size performs well in finite-horizon markets. A planner can implement this simple heuristic with almost no knowledge of the market; the planner only needs to know the market's time horizon to select an appropriate adjustment size.

This simple SGD heuristic gives rise to a form of price rigidity. The optimal adjustment size balances two potential losses. A larger adjustment size facilitates a faster adjustment toward market-clearing prices, but a smaller adjustment size creates rigidity that protects from overreaction to noise. Intuitively, the optimal adjustment size balances the two, and induces more rigidity when the time horizon T indicates a more stable environment. Because of this trade-off, the SGD heuristic can appear to be too rigid and slow to adjust prices in initial periods.

References

Adan, I., Bušić, A., Mairesse, J., and Weiss, G. (2018). Reversibility and further properties of fcfs infinite bipartite matching. *Mathematics of Operations Research*, 43(2):598–621.

Adan, I. and Weiss, G. (2012). Exact fcfs matching rates for two infinite multitype sequences. *Operations research*, 60(2):475–489.

Agarwal, N., Ashlagi, I., Rees, M. A., Somaini, P., and Waldinger, D. (2021). Equilibrium allocations under alternative waitlist designs: Evidence from deceased donor kidneys. *Econometrica*, 89(1):37–76.

Agarwal, N., Hodgson, C., and Somaini, P. (2020). Choices and outcomes in assignment mechanisms: The allocation of deceased donor kidneys. Technical report, National Bureau of Economic Research.

Akbarpour, M., Dworczak, P., and Kominers, S. D. (2020). Redistributive allocation mechanisms. *Available at SSRN 3609182*.

Arnosti, N. and Shi, P. (2020). Design of lotteries and wait-lists for affordable housing allocation. *Management Science*.

Baccara, M., Lee, S., and Yariv, L. (2018). Optimal dynamic matching. Technical report, CEPR Discussion Papers.

Bertsekas, D. P. (1981). A new algorithm for the assignment problem. *Mathematical Programming*, 21(1):152–171.

Bertsimas, D., Gamarnik, D., and Tsitsiklis, J. N. (2001). Performance of multiclass markovian queueing networks via piecewise linear lyapunov functions. *Annals of Applied Probability*, pages 1384–1428.

Bloch, F. and Cantala, D. (2017). Dynamic assignment of objects to queuing agents. *American Economic Journal: Microeconomics*, 9(1):88–122.

Bubeck, S., Devanur, N. R., Huang, Z., and Niazadeh, R. (2019). Multi-scale online learning: Theory and applications to online auctions and pricing. *The Journal of Machine Learning Research*, 20(1):2248–2284.

Buchholz, N., Doval, L., Kastl, J., Matějka, F., and Salz, T. (2020). The value of time: Evidence from auctioned cab rides. Technical report, National Bureau of Economic Research.

Caldentey, R., Kaplan, E. H., and Weiss, G. (2009). Fcfs infinite bipartite matching of servers and customers. *Advances in Applied Probability*, 41(3):695–730.

Che, Y.-K. and Tercieux, O. (2021). Optimal queue design. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 312–313.

Cheung, Y. K., Cole, R., and Devanur, N. R. (2019). Tatonnement beyond gross substitutes? gradient descent to the rescue. *Games and Economic Behavior*.

Cheung, Y. K., Cole, R., and Tao, Y. (2018). Dynamics of distributed updating in fisher markets. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 351–368.

Cole, R. and Fleischer, L. (2008). Fast-converging tatonnement algorithms for one-time and ongoing market problems. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 315–324.

Condorelli, D. (2013). Market and non-market mechanisms for the optimal allocation of scarce resources. *Games and Economic Behavior*, 82:582–591.

Correa, J. R. and Stier-Moses, N. E. (2010). Wardrop equilibria. *Wiley encyclopedia of operations research and management science.*

Demange, G., Gale, D., and Sotomayor, M. (1986). Multi-item auctions. *Journal of Political Economy*, 94(4):863–872.

Deng, S. (1998). Global error bounds for convex inequality systems in banach spaces. *SIAM journal on control and optimization*, 36(4):1240–1249.

Devanur, N. R., Jain, K., Sivan, B., and Wilkens, C. A. (2019). Near optimal online algorithms and fast approximation algorithms for resource allocation problems. *Journal of the ACM (JACM)*, 66(1):1–41.

Dworczak, P., Kominers, S. D., and Akbarpour, M. (2021). Redistribution through markets. *Econometrica*, 89(4):1665–1698.

Galichon, A. (2018). *Optimal transport methods in economics.* Princeton University Press.

Gallier, J. (2008). Notes on convex sets, polytopes, polyhedra, combinatorial topology, voronoi diagrams and delaunay triangulations. *arXiv preprint arXiv:0805.0292.*

Hartline, J. D. and Roughgarden, T. (2008). Optimal mechanism design and money burning. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 75–84.

Hassin, R. (2016). *Rational queueing.* CRC press.

Hassin, R. and Haviv, M. (2003). *To queue or not to queue: Equilibrium behavior in queueing systems*, volume 59. Springer Science & Business Media.

Johari, R., Kamble, V., and Kanoria, Y. (2016). Matching while learning. *arXiv preprint arXiv:1603.04549.*

Kaplan, E. H. (1984). *Managing the demand for public housing.* PhD thesis, Massachusetts Institute of Technology.

Kaplan, E. H. (1988). A public housing queue with reneging. *Decision Sciences*, 19(2):383–391.

Kelly, F. and Yudovina, E. (2018). A markov model of a limit order book: thresholds, recurrence, and trading strategies. *Mathematics of Operations Research*, 43(1):181–203.

Kozakiewicz, W. et al. (1947). On the convergence of sequences of moment generating functions. *The Annals of Mathematical Statistics*, 18(1):61–69.

LeCun, Y. A., Bottou, L., Orr, G. B., and Müller, K.-R. (2012). Efficient backprop. In *Neural networks: Tricks of the trade*, pages 9–48. Springer.

Leshno, J. (2017). Dynamic matching in overloaded waiting lists. *Available at SSRN 2967011.*

Loertscher, S., Muir, E. V., and Taylor, P. G. (2018). Optimal market thickness and clearing. *Unpublished paper, Department of Economics, University of Melbourne.[1224].*

Mendelson, H. (1982). Market behavior in a clearing house. *Econometrica: Journal of the Econometric Society*, pages 1505–1524.

Myerson, R. B. (1981). Optimal auction design. *Mathematics of operations research*, 6(1):58–73.

Naor, P. (1969). The regulation of queue size by levying tolls. *Econometrica: journal of the Econometric Society*, pages 15–24.

Postnikov, A. (2009). Permutohedra, associahedra, and beyond. *International Mathematics Research Notices*, 2009(6):1026–1106.

Powell, W. B. and Sheffi, Y. (1982). The convergence of equilibrium algorithms with predetermined step sizes. *Transportation Science*, 16(1):45–55.

Roth, A., Slivkins, A., Ullman, J., and Wu, Z. S. (2020). Multidimensional dynamic pricing for welfare maximization. *ACM Transactions on Economics and Computation (TEAC)*, 8(1):1–35.

Su, X. and Zenios, S. (2004). Patient choice in kidney allocation: The role of the queueing discipline. *Manufacturing & Service Operations Management*, 6(4):280–301.

Uzawa, H. (1960). Walras' tatonnement in the theory of exchange. *The Review of Economic Studies*, 27(3):182–194.

van Dijk, W. (2019). The socio-economic consequences of housing assistance. *University of Chicago Kenneth C. Griffin Department of Economics job market paper, 0–46 i–xi*, 36.

Verdier, V. and Reeling, C. (2022). Welfare effects of dynamic matching: An empirical analysis. *The Review of Economic Studies*, 89(2):1008–1037.

Waldinger, D. (2021). Targeting in-kind transfers through market design: A revealed preference analysis of public housing allocation. *American Economic Review*, 111(8):2660–96.

Weiss, G. (2021). *Scheduling and Control of Queueing Networks*, volume 14. Cambridge University Press.

Weitzman, M. L. (1977). Is the price system or rationing more effective in getting a commodity to those who need it most? *The Bell Journal of Economics*, pages 517–524.

Wolff, R. W. (1982). Poisson arrivals see time averages. *Operations Research*, 30(2):223–231.

Zenios, S. A. (1999). Modeling the transplant waiting list: A queueing model with reneging. *Queueing systems*, 31(3-4):239–251.

A Proof of Theorem 1 for General Waiting Costs

In this section we complete the proof of Theorem 1 for the general case. The proof is identical to the proof in the main body except for the second step of Lemma 4. The following three lemmas establish inequalities that give versions of inequalities (11) and (12) for general waiting cost functions that satisfy the conditions in Section 2. Replacing Step 2 in the proof of Lemma 4 with these lemmas gives the general proof.

We use the same Lyapunov function in the analysis, where $L(\mathbf{q})$ is such that $\nabla L(\mathbf{q}) = \mathbf{p}(\mathbf{q})$. The analysis uses the Bregman divergence generated by $L(\mathbf{q})$ as the notion of proximity, which is defined as

$$D_L(\mathbf{q}_1, \mathbf{q}_2) \triangleq L(\mathbf{q}_1) - L(\mathbf{q}_2) - \langle \nabla L(\mathbf{q}_2), \mathbf{q}_1 - \mathbf{q}_2 \rangle \ .$$

Lemma 5. *We have*

$$L(\mathbf{q}_{t+1}) \leq L(\mathbf{q}_t + \mathbf{a}_t - \mathbf{g}_t) + \frac{\Delta}{2} \cdot \mathbb{1}_{\{\xi_t = 0\}} \, .$$

Proof. By definition of Bregman divergence, we have

$$\begin{aligned} D_L(\mathbf{q}_t + \mathbf{a}_t - \mathbf{g}_t, \mathbf{q}_{t+1}) &= L(\mathbf{q}_t + \mathbf{a}_t - \mathbf{g}_t) - L(\mathbf{q}_{t+1}) - \langle \mathbf{p}(\mathbf{q}_{t+1}), \mathbf{q}_t + \mathbf{a}_t - \mathbf{g}_t - \mathbf{q}_{t+1} \rangle \\ &= L(\mathbf{q}_t + \mathbf{a}_t - \mathbf{g}_t) - L(\mathbf{q}_{t+1}) + \langle \mathbf{p}(\mathbf{q}_{t+1}), \mathbf{u}_t \rangle \, . \end{aligned}$$

Therefore,

$$L(\mathbf{q}_{t+1}) = \ L(\mathbf{q}_t + \mathbf{a}_t - \mathbf{g}_t) + \langle \mathbf{p}(\mathbf{q}_{t+1}), \mathbf{u}_t \rangle - D_L(\mathbf{q}_t + \mathbf{a}_t - \mathbf{g}_t, \mathbf{q}_{t+1}) \, . \tag{20}$$

To bound the RHS of (20), we consider two cases:

Case 1. If $\exists j \in \mathcal{J}$ such that $d_{j,t} = 1$ and $q_{j,t} = 0$, we have $q_{j,t+1} = 0$ and $l_{j,t} = 1$. Note that in this case $\xi_t = 0$. Let $P_j(q)$ be an anti-derivative of $p_j(q)$; then $L(\mathbf{q}) = \sum_{j \in \mathcal{J}} P_j(q)$ is a Lyapunov function because it satisfies $\nabla L(\mathbf{q}) = \mathbf{p}(\mathbf{q})$. We have

$$\begin{aligned} &\langle \mathbf{p}(\mathbf{q}_{t+1}), \mathbf{u}_t \rangle - D_L(\mathbf{q}_t + \mathbf{a}_t - \mathbf{g}_t, \mathbf{q}_{t+1}) \\ &= p_j(0) - (P_j(-1) - P_j(0) - p_j(0) \cdot (-1)) \\ &= P_j(0) - P_j(-1) \, . \end{aligned} \tag{21}$$

Since $p_j(\cdot)$ is nonnegative and Δ-Lipshitz, we have

$$P_j(0) \leq P_j(-1) + \int_0^1 \Delta \cdot x dx = P_j(-1) + \Delta/2 \, .$$

Plugging the above equality into (21), we have

$$\langle \mathbf{p}(\mathbf{q}_{t+1}), \mathbf{u}_t \rangle - D_L(\mathbf{q}_t + \mathbf{a}_t - \mathbf{g}_t, \mathbf{q}_{t+1}) \le \frac{\Delta}{2} \cdot \mathbb{1}_{\{\xi_t = 0\}} . \tag{22}$$

Case 2. If the condition in Case 1 does not hold, we have $\mathbf{u}_t = \mathbf{0}$ and $\mathbf{q}_t + \mathbf{a}_t - \mathbf{g}_t = \mathbf{q}_{t+1}$; hence

$$\langle \mathbf{p}(\mathbf{q}_{t+1}), \mathbf{u}_t \rangle - D_L(\mathbf{q}_t + \mathbf{a}_t - \mathbf{g}_t, \mathbf{q}_{t+1}) = 0 .$$

Therefore, plugging the above two cases into (20), we have

$$L(\mathbf{q}_{t+1}) \le L(\mathbf{q}_t + \mathbf{a}_t - \mathbf{g}_t) + \frac{\Delta}{2} \cdot \mathbb{1}_{\{\xi_t = 0\}} .$$

□

Lemma 6. *We have*

$$\mathbb{E}[v(\theta_t, a(\theta_t, \mathbf{q}_t)) | \mathbf{q}_t] \ge \frac{\lambda}{1+\lambda} h(\mathbf{p}_t) - (L(\mathbf{q}_t) - \mathbb{E}[L(\mathbf{q}_{t+1}) \mid \mathbf{q}_t]) - \frac{\Delta}{2(1+\lambda)} - \mathbb{E}[D_L(\mathbf{q}_t + \mathbf{a}_t - \mathbf{g}_t, \mathbf{q}_t) \mid \mathbf{q}_t] .$$

Proof. We have that the drift of Lyapunov function $L(\mathbf{q})$ in one period is

$$\begin{aligned} & L(\mathbf{q}_t) - L(\mathbf{q}_{t+1}) \\ \ge\ & L(\mathbf{q}_t) - L(\mathbf{q}_t + \mathbf{a}_t - \mathbf{g}_t) - \frac{\Delta}{2} \cdot \mathbb{1}_{\{\xi_t = 0\}} \\ =\ & -\langle \mathbf{p}(\mathbf{q}_t), \mathbf{a}_t - \mathbf{g}_t \rangle - D_L(\mathbf{q}_t + \mathbf{a}_t - \mathbf{g}_t, \mathbf{q}_t) - \frac{\Delta}{2} \cdot \mathbb{1}_{\{\xi_t = 0\}} , \end{aligned} \tag{23}$$

where the inequality follows from Lemma 5, and the equality follows from the definition of Bregman divergence.

Adding $v(\theta_t, a(\theta_t, \mathbf{q}_t))$ to both sides of equation (23), we have

$$\begin{aligned} & v(\theta_t, a(\theta_t, \mathbf{q}_t)) + L(\mathbf{q}_t) - L(\mathbf{q}_{t+1}) \\ \ge\ & v(\theta_t, a(\theta_t, \mathbf{q}_t)) - \langle \mathbf{p}(\mathbf{q}_t), \mathbf{a}_t - \mathbf{g}_t \rangle - D_L(\mathbf{q}_t + \mathbf{a}_t - \mathbf{g}_t, \mathbf{q}_t) - \frac{\Delta}{2} \cdot \mathbb{1}_{\{\xi_t = 0\}} . \end{aligned}$$

Taking expectation conditional on $\mathbf{q}_t$ and applying equation (10), we have

$$\begin{aligned} & \mathbb{E}[v(\theta_t, a(\theta_t, \mathbf{q}_t)) | \mathbf{q}_t] + (L(\mathbf{q}_t) - \mathbb{E}[L(\mathbf{q}_{t+1}) \mid \mathbf{q}_t]) \\ \ge\ & \frac{\lambda}{1+\lambda} h(\mathbf{p}_t) - \mathbb{E}[D_L(\mathbf{q}_t + \mathbf{a}_t - \mathbf{g}_t, \mathbf{q}_t) \mid \mathbf{q}_t] - \frac{\Delta}{2(1+\lambda)} . \end{aligned}$$

Rearranging the terms, we obtain the desired inequality. □

Lemma 7. *We have that for any* $\mathbf{q}_t$*,*

$$\mathbb{E}[D_L(\mathbf{q}_t + \mathbf{a}_t - \mathbf{g}_t, \mathbf{q}_t) \mid \mathbf{q}_t] \leq \frac{\Delta}{2} \,.$$

Proof. Note that $L(\mathbf{q})$ is convex because its gradient $\nabla L(\mathbf{q}) = \mathbf{p}(\mathbf{q})$ is increasing in each coordinate. Also note that $L(\mathbf{q})$ has a Δ-Lipschitz gradient, because for queue lengths $\mathbf{q}_1, \mathbf{q}_2$,

$$||\nabla L(\mathbf{q}_1) - \nabla L(\mathbf{q}_2)|| = ||\mathbf{p}(\mathbf{q}_1) - \mathbf{p}(\mathbf{q}_2)|| \leq \Delta ||\mathbf{q}_1 - \mathbf{q}_2|| \,.$$

Equivalently, $L(\mathbf{q})$ is Δ-strongly smooth, i.e.,

$$L(\mathbf{q}_2) - L(\mathbf{q}_1) \leq \langle \nabla L(\mathbf{q}_1), \mathbf{q}_2 - \mathbf{q}_1 \rangle + \frac{\Delta}{2} ||\mathbf{p}(\mathbf{q}_2) - \mathbf{p}(\mathbf{q}_1)||^2 \,.$$

By the definition of Bregman divergence, we have

$$\begin{aligned} & D_L(\mathbf{q}_t + \mathbf{a}_t - \mathbf{g}_t, \mathbf{q}_t) \\ &= L(\mathbf{q}_t + \mathbf{a}_t - \mathbf{g}_t) - L(\mathbf{q}_t) - \langle \nabla L(\mathbf{q}_t), \mathbf{a}_t - \mathbf{g}_t \rangle \\ &\leq \frac{\Delta}{2} ||\mathbf{a}_t - \mathbf{g}_t||^2 \\ &= \frac{\Delta}{2} \,, \end{aligned}$$

where the inequality follows from the strong smoothness of $L(\mathbf{q})$. This completes the proof. □

B Omitted Proofs

B.1 Proof of Proposition 1 from Section 2

In this section we prove Proposition 1, showing that $W^{\text{OPT}} = W^*$.

Proof of Proposition 1. We first show that $W^{\text{OPT}} \leq W^*$. This part of the proof mostly consists of a careful treatment of expectations and limits.

Let $\eta \in H$ be any no-Ponzi allocation. Without loss of generality, assume that $\Theta = [0, 1]$ and that F is a CDF on $[0, 1]$. This allows us to rewrite problem (3) as

$$W^* = \max_{G_j(\theta),\, j \in \mathcal{J}} \quad \sum_{j \in \mathcal{J}} \int_\Theta v(\theta, j) \, dG_j(\theta)$$

$$\begin{aligned} \text{subject to} \quad & \sum_{j \in \mathcal{J}} G_j(\theta) \leq F(\theta)\,, && \forall \theta\,, \\ & G_j(0) = 0;\; G_j(1) \leq \mu_j/\lambda\,, && \forall j\,, \\ & G_j(\cdot) \text{ is nondecreasing and right-continuous}\,, && \forall j\,. \end{aligned} \tag{24}$$

Let $A_T = |A(T)|$, and recall that $A(T)$ is the set of epochs up to T in which an agent arrived. Recall that $\eta_t \in \mathcal{J}_\emptyset$ is the kind of item assigned under η to the agent that arrived in epoch t. For each $j \in \mathcal{J}$ and $\theta \in \Theta$, define $\hat{G}_j^T(\theta)$ as

$$\hat{G}_j^T(\theta) \triangleq \frac{1}{A_T} \sum_{t \leq T} \mathbb{1}_{\{\xi_t = 1, \theta_t \leq \theta, \eta_t = j\}}\,.$$

Therefore, $\hat{G}_j^T(\theta)$ is proportional to the empirical cumulative distribution function of the types of the agents who arrived in epochs in $A(T)$ and are assigned a type j item. When $A_T = 0$, we set $\hat{G}_j^T(\theta) = 0$ for all $j \in \mathcal{J}$ and $\theta \in \Theta$. By definition, the allocative efficiency under η is defined as

$$W(\eta) = \liminf_{T \to \infty} \frac{1}{A_T} \sum_{j \in \mathcal{J}} \sum_{t \leq T} \xi_t v(\theta_t, j) = \liminf_{T \to \infty} \sum_{j \in \mathcal{J}} \int_{\theta \in \Theta} v(\theta, j) d\hat{G}_j^T(\theta)\,. \tag{25}$$

Note that for any T, $\hat{G}_j^T(\theta)$ satisfies

$$\sum_{j \in \mathcal{J}} \hat{G}_j^T(\theta) \leq \frac{1}{A_T} \sum_{t \leq T} \xi_t \mathbb{1}_{\{\theta_t \leq \theta\}}\,, \qquad \forall \theta \in \Theta\,, \tag{26}$$

$$\hat{G}_j^T(0) = 0\,, \hat{G}_j^T(1) \leq \frac{1}{A_T} \left(\sum_{t \leq T} \mathbb{1}_{\{\eta_t = j\}} + M \right), \qquad \forall j \in \mathcal{J}\,, \tag{27}$$

$$\hat{G}_j^T(\theta) \text{ is nondecreasing and right-continuous}, \qquad \forall j \in \mathcal{J}\,, \tag{28}$$

for some $M \in \mathbb{R}$. Here (26) and (28) are trivial. Equality (27) is satisfied by any no-Ponzi assignment for the following reason: the agents arriving in epochs $A(T)$ who are assigned a type $j \in \mathcal{J}$ item are either assigned before the T-th epoch or after the T-th epoch. The number of those who are assigned before T cannot exceed the total number of type j items that arrive before T. The number of those who are assigned after T is bounded by some $M \in \mathbb{R}$ by the definition of no-Ponzi assignments.

Combining the above, we have

$$\mathbb{E}\left[\sum_{j \in \mathcal{J}} \int_{\theta \in \Theta} v(\theta, j) d\hat{G}_j^T(\theta) \right] \leq \mathbb{E}\left[\max_{\hat{G}_j(\theta) \text{ satisfying } (26)(27)(28)} \sum_{j \in \mathcal{J}} \int_{\theta \in \Theta} v(\theta, j) d\hat{G}_j(\theta) \right].$$

It is easy to check that the optimal value of the inner maximization problem above is concave and nondecreasing in the RHS of (26) and (27). Note that

$$\begin{aligned}
&\text{Expectation of RHS of (26)} = F(\theta)\,, && \forall \theta \in \Theta\,,\\
&\text{Expectation of RHS of (27)} = \frac{\mu_j}{\lambda} + \frac{(1+\lambda)M}{\lambda T}\,, && \forall j \in \mathcal{J}\,.
\end{aligned}$$

It follows from Fatou's Lemma that

$$\mathbb{E}[W(\eta)] = \mathbb{E}\left[\liminf_{T\to\infty} \sum_{j\in\mathcal{J}} \int_{\theta\in\Theta} v(\theta,j)d\hat{G}_j^T(\theta)\right] \le \liminf_{T\to\infty} \mathbb{E}\left[\sum_{j\in\mathcal{J}} \int_{\theta\in\Theta} v(\theta,j)d\hat{G}_j^T(\theta)\right].$$

Applying Jensen's inequality, we have

$$\liminf_{T\to\infty} \mathbb{E}\left[\sum_{j\in\mathcal{J}} \int_{\theta\in\Theta} v(\theta,j)d\hat{G}_j^T(\theta)\right] \le W^*\,,$$

where W^* is defined in (3). Therefore $\mathbb{E}[W(\eta)] \le W^*$ for any $\eta \in H$. Since $W(\eta)$ is uniformly bounded above by $v_{\max}$, by the Bounded Convergence Theorem we have $W^{\mathrm{OPT}} = \mathbb{E}[\sup_{\eta\in H} W(\eta)] = \sup_{\eta\in H} \mathbb{E}[W(\eta)] \le W^*$. This completes the first part of the proof.

Next we prove that $W^{\mathrm{OPT}} \ge W^*$. We explicitly construct a sequence of randomized policies that can achieve allocative efficiencies that are arbitrarily close to W^*. Note that the constructed policies are technical devices used to prove the desired bound, rather than practical policies.

Denote the optimal solution of the optimization problem (3) by $\mathbf{x}^*$. Consider the following randomized policy: maintain a separate First-Come First-Served queue for each item. An arriving agent will be assigned to one of the queues or rejected based on a coin toss (to be specified later). An agent who joins a queue will wait in that queue until receiving an item. When an item arrives, it is assigned to the agent at the head of its queue; if the queue is empty, the item is discarded. The coin toss is defined as follows: fix $M \in \mathbb{Z}_+$. If the arriving agent is of type θ, she is assigned to a queue j with probability $x^*_{\theta j}$, or rejected from all queues with probability $1-\sum_{j\in\mathcal{J}} x^*_{\theta j}$. If the length of the queue to which the agent is assigned exceeds M, the agent is also rejected.

Denote the match value collected by the randomized policy in epoch t by v_t^{RD}. Then by definition of the policy, we have

$$\mathbb{E}[v_t^{\mathrm{RD}} | q_{j,t} < M, \forall j \in \mathcal{J}] = \frac{\lambda}{1+\lambda} \sum_{j\in\mathcal{J}} \int_{\Theta} x^*_{\theta j}\, v(\theta,j)dF(\theta) = \frac{\lambda}{1+\lambda} W^*\,.$$

It follows that

$$\mathbb{E}[v_t^{\mathrm{RD}}] \geq \frac{\lambda}{1+\lambda} W^* \cdot \mathbb{P}(q_{j,t} < M, \forall j \in \mathcal{J})\,.$$

Let W^{RD} be the allocative efficiency of the randomized policy. We therefore have

$$W^{\mathrm{RD}} \geq W^* \cdot \mathbb{P}(q_{j,\infty} < M, \forall j \in \mathcal{J})\,,$$

where $\mathbf{q}_\infty$ is the steady-state queue length distribution. The allocative efficiency loss of the randomized policy can be bounded by

$$\begin{aligned} W^* - W^{\mathrm{RD}} &\leq W^* - W^* \cdot \mathbb{P}(q_{j,\infty} < M, \forall j \in \mathcal{J}) \\ &\leq v_{\max} \cdot \mathbb{P}(q_{j,\infty} = M, \exists j \in \mathcal{J}) \\ &\leq v_{\max} \cdot \sum_{j \in \mathcal{J}} \mathbb{P}(q_{j,\infty} = M)\,, \end{aligned} \tag{29}$$

where the second inequality follows from the fact that $W^* \leq v_{\max}$, and the last inequality comes from the union bound. It remains to bound $\mathbb{P}(q_{j,\infty} = M)$ for any $j \in \mathcal{J}$ under the randomized policy. Fix $j \in \mathcal{J}$; then $q_{j,t}$ is a birth–death process on $\{0, 1, \cdots, M\}$ with death rate μ_j and birth rate

$$\lambda_j^* = \lambda \int_\Theta x_{\theta j}^* dF(\theta)\,.$$

It follows from the constraint in (3) that $\lambda_j^* \leq \mu_j$. As a result, $\mathbb{P}(q_{j,\infty} = M) \leq \mathbb{P}(q_{j,\infty} = M-1) \leq \cdots \leq \mathbb{P}(q_{j,\infty} = 0)$; hence $\mathbb{P}(q_{j,\infty} = M) \leq \frac{1}{M+1}$. Plugging in the bound on $\mathbb{P}(q_{j,\infty} = M)$ to (29), we have

$$W^* - W^{\mathrm{RD}} \leq \frac{v_{\max}|\mathcal{J}|}{M+1}\,.$$

Notice that by definition, $W^{\mathrm{RD}} \leq W^{\mathrm{OPT}}$, and hence

$$W^* - W^{\mathrm{OPT}} \leq \frac{v_{\max}|\mathcal{J}|}{M+1}\,.$$

Since M can be chosen arbitrarily, it must be true that $W^* - W^{\mathrm{OPT}} \leq 0$. This completes the proof. □

B.2 Omitted proofs from Section 3

The following lemma was used in the proof of Theorem 1. Its proof uses standard arguments, and is presented here for completeness.

Lemma 8. *The following equality holds:*

$$W^{\mathrm{WL}} = \frac{1+\lambda}{\lambda} \lim_{T\to\infty} \frac{1}{T}\mathbb{E}\left[W_T(\eta^{\mathrm{WL}})\right] . \tag{30}$$

Proof. We need to argue that the limiting operator and the expectation can be interchanged. This is done in the following two steps.

First, we show that the stochastic process $\{(\xi_t, v(\theta_t, a(\theta_t, \mathbf{q}_t)))\}_{t\geq 0}$ is ergodic. Note that $v(\theta_t, a(\theta_t, \mathbf{q}_t))$ only depends on $\mathbf{q}_t$ and independent variables θ_t. The finite-state Markov chain $\{\mathbf{q}_t\}_{t\geq 0}$ is irreducible and aperiodic;[28] therefore it has a unique steady-state distribution and $\{(\xi_t, v(\theta_t, a(\theta_t, \mathbf{q}_t)))\}_{t\geq 0}$ is ergodic.

Second, we exchange the order of the limiting operator and the expectation. It follows from Birkhoff's ergodic theorem that $\frac{W_T(\eta^{\mathrm{WL}})}{A_T}$ converges almost surely to $\mathbb{E}[v_\infty|\xi_\infty = 1]$, where (ξ_∞, v_∞) is the steady-state distribution of $(\xi_t, v(\theta_t, a(\theta_t, \mathbf{q}_t)))$. Since $\frac{W_T(\eta^{\mathrm{WL}})}{A_T}$ is nonnegative and uniformly bounded from above by $v_{\max}$ for all $T > 0$, we have

$$W^{\mathrm{WL}} = \mathbb{E}\left[\lim_{T\to\infty} \frac{W_T(\eta^{\mathrm{WL}})}{A_T}\right] = \lim_{T\to\infty} \mathbb{E}\left[\frac{W_T(\eta^{\mathrm{WL}})}{A_T}\right] = \mathbb{E}[v_\infty|\xi_\infty = 1] ,$$

where we apply the Bounded Convergence Theorem in the second equality to exchange the limits; the last equality holds because the boundedness of $\frac{W_T(\eta^{\mathrm{WL}})}{A_T}$ and its almost sure convergence imply L_1 convergence. Finally, observe that

$$\begin{aligned}
\mathbb{E}[v_\infty] &= \mathbb{E}[v_\infty|\xi_\infty = 1]\cdot \mathbb{P}(\xi_\infty = 1) + \mathbb{E}[v_\infty|\xi_\infty = 0]\cdot \mathbb{P}(\xi_\infty = 0) \\
&= \mathbb{E}[v_\infty|\xi_\infty = 1]\cdot \frac{\lambda}{1+\lambda} + 0 \\
&= W^{\mathrm{WL}} \cdot \frac{\lambda}{1+\lambda} ,
\end{aligned}$$

where the second equality follows from the fact that all the rewards are collected when agents arrive, i.e., $\xi_t = 1$. Note that

$$\lim_{T\to\infty} \frac{1}{T}\mathbb{E}\left[W_T(\eta^{\mathrm{WL}})\right] = \mathbb{E}[v_\infty] ,$$

[28] Irreducibility follows from the fact that all states can go to $\mathbf{0}$ with positive probability. Aperiodicity comes from the fact that the state can stay at $\mathbf{0}$ for an arbitrary number of periods.

and therefore we have

$$W^{\mathrm{WL}} = \frac{1+\lambda}{\lambda}\mathbb{E}[v_\infty] = \frac{1+\lambda}{\lambda}\lim_{T\to\infty}\frac{1}{T}\mathbb{E}\left[W_T(\eta^{\mathrm{WL}})\right] .$$

□

Proof of Theorem 2. The proof is identical to the proof of Theorem 1 presented in Section 4, except for the following changes. We replace inequality (12) with the following inequality,

$$\frac{1}{2}\sum_{j\in\mathcal{J}}\frac{c}{\mu_j}\mathbb{E}[(a_{j,t}-g_{j,t})^2 + l_{j,t}^2 \mid \mathbf{q}_t] \le \frac{1}{2}\Delta + \varepsilon\,, \tag{31}$$

and show there exist $\alpha, \beta, c_0 > 0$, such that for any $c_\ell < c_0$ we have $\varepsilon < \beta \cdot e^{-\alpha/\Delta_\ell}$.

To show inequality (31), observe that

$$\frac{c}{2\mu_j}(a_{j,t}-g_{j,t})^2 + \frac{c}{2\mu_j}l_{j,t}^2 = \begin{cases} \frac{c}{\mu_j} & \text{if } g_{j,t}=1 \text{ and } q_j = 0\,, \\ \frac{c}{2\mu_j} & \text{otherwise.} \end{cases}$$

That is, the above term is equal to $c/\mu_j \le \Delta$ for arrivals that correspond to an item that is discarded because its queue is empty, and is equal to $c/2\mu_j \le \Delta/2$ for all other arrivals. By Corollary 3, there exist $\alpha, \beta, c_0 > 0$, such that for every market with $c_\ell < c_0$ the probability that an arriving item j is discarded is at most $\beta \cdot e^{-\alpha/\Delta_\ell}$. Thus,

$$\begin{aligned}&\frac{1}{2}\sum_{j\in\mathcal{J}}\frac{c}{\mu_j}\mathbb{E}[(a_{j,t}-g_{j,t})^2 + l_{j,t}^2 \mid \mathbf{q}_t] \\ &\le \frac{\Delta}{2} + \beta\cdot e^{-\alpha/\Delta_\ell}\frac{\Delta}{2}\,.\end{aligned}$$

Using (31), we can replace (8) in the proof of Theorem 1 with

$$\begin{aligned}\mathbb{E}\left[W_T(\eta^{\mathrm{WL}})\right] &= \mathbb{E}\left[\sum_{t=1}^{T}\xi_t\cdot v(\theta_t, a(\theta_t,\mathbf{q}_t))\right] \\ &\ge \mathbb{E}\left[\sum_{t=1}^{T}\frac{\lambda}{1+\lambda}W^* - \frac{1}{\mu_{\min}\cdot\Delta}\left(L(\mathbf{p}_t) - \mathbb{E}[L(\mathbf{p}_{t+1}\mid \mathbf{q}_t])\right) - \frac{1}{2}\Delta - \varepsilon\right] \\ &= T\frac{\lambda}{1+\lambda}W^* - \frac{1}{\mu_{\min}\cdot\Delta}\left(L(\mathbf{p}_1) - \mathbb{E}[L(\mathbf{p}_{T+1})]\right) - T\frac{1}{2}\Delta - T\varepsilon\,. \end{aligned} \tag{32}$$

Following the same steps as in the proof of Theorem 1 and using the assumption that

$\lambda \geq 1$, we have

$$\begin{aligned} W^{\mathrm{WL}} &= \frac{1+\lambda}{\lambda} \lim_{T\to\infty} \frac{1}{T}\mathbb{E}\left[W_T(\eta^{\mathrm{WL}})\right] \\ &\geq W^* - \frac{1+\lambda}{2\lambda}\Delta - \frac{1+\lambda}{\lambda}\varepsilon \\ &\geq W^* - \Delta - 2\varepsilon\,, \end{aligned}$$

which completes the proof.

□

B.3 Omitted proofs from Section 3.2

Proof of Proposition 2. We prove the result by calculating the allocative efficiency loss in the market of Example 2. By Proposition 1, we have that $W^{\mathrm{OPT}} = \gamma$.

Under the waiting list, an agent of type θ will only join the queue for item $j = \theta$. A type $\theta_t = j$ agent arriving in epoch t will join queue j to receive a value of γ only if

$$\gamma \geq p_j(q_j) = \frac{c}{\mu_j}(1 + q_{t,j}) = \Delta(1 + q_{t,j})\ ,$$

or

$$q_{t,j} \leq \frac{\gamma}{\Delta} - 1\ .$$

Therefore, the possible states of each queue $j \in \mathcal{J}$ are $0, 1, \ldots, K$ with $K = \lfloor \gamma/\Delta \rfloor$. Let $\pi_j(k)_{0\leq k\leq K}$ denote the steady-state distributions over the length of queue j. Because the length of the queue follows a reflected unbiased random walk, all states are equally likely and[29] $\pi_j(k) = \frac{1}{K+1}$.

The allocative efficiency loss under the waiting list is given by

$$\begin{aligned} W^{\mathrm{OPT}} - W^{\mathrm{WL}} &= \gamma - \sum_{j\in\Theta} F(j) \left(\pi_j(K)\cdot 0 + \sum_{k<K} \pi_j(k)\cdot\gamma\right) \\ &= \gamma - J\frac{1}{J}\left(\frac{1}{K+1}\cdot 0 + \frac{K}{K+1}\cdot\gamma\right) \\ &= \frac{1}{K+1}\cdot\gamma \\ &= \frac{1}{\lfloor \gamma/\Delta \rfloor + 1}\cdot\gamma,. \end{aligned}$$

[29]To see this directly, observe that equating probability flows across a cut gives for any $0 < k < K$ that $\pi_j(k)\lambda/J = \pi_j(k+1)\mu_j$, which implies that $\pi_j(0) = \pi_j(1) = \cdots = \pi_j(K)$.

By choosing γ such that $\lfloor \gamma/\Delta \rfloor \approx \gamma/\Delta - 1$ we get that

$$\begin{aligned} W^{\mathrm{OPT}} - W^{\mathrm{WL}} &\geq \frac{1}{\gamma/\Delta - 1 + 1} \cdot \gamma - \varepsilon \\ &= \Delta - \varepsilon \, . \end{aligned}$$

□

B.4 Omitted proofs from Section 3.3

Proof of Lemma 1. A problem instance $(\boldsymbol{\lambda}, \boldsymbol{\mu})$ is said to satisfy *generalized imbalance (GI)* if there is no pair of nonempty subsets of agent types $\mathcal{I}' \subset \Theta$ and item types $\mathcal{J}' \subset \mathcal{J}$ such that the total arrival rate of agents with type in $\mathcal{I}'$ exactly matches the total arrival rate of items with types in $\mathcal{J}'$, that is,

$$\sum_{j \in \mathcal{J}'} \mu_j \neq \sum_{\theta \in \mathcal{I}'} \lambda_\theta \qquad \forall \mathcal{J}' \subset \mathcal{J},\, \mathcal{I}' \subset \Theta \, .$$

The set of $(\boldsymbol{\lambda}, \boldsymbol{\mu})$ satisfying GI is open and dense in $\mathbb{R}_{++}^{|\mathcal{I}|} \times \mathbb{R}_{++}^{|\mathcal{J}|}$. The result follows because GI implies uniqueness of market-clearing prices (see Proposition C.2 of Johari et al., 2016).

□

B.5 Omitted proofs from Section 3.4

Proof of Proposition 3. For ease of reading, the proof is divided into three steps. First, we provide a lower bound for the dual objective. Second, we use that lower bound to construct a lower bound on the Lyapunov drift. Third, we derive a concentration bound using the Lyapunov drift.

Step 1. We first establish a lower bound of the dual function $h(\mathbf{p})$; namely, for any $\epsilon > 0$, there exists $\gamma > 0$, such that $\forall \mathbf{p} \notin B_\epsilon(\mathcal{P}^*)$,

$$h(\mathbf{p}) - h(\mathbf{p}^*) \geq \gamma \cdot \mathrm{dist}(\mathbf{p}, \mathcal{P}^*) \, . \tag{33}$$

To establish (33), we use the ϵ-level set of the dual function $h(\mathbf{p})$, which we denote by $\mathcal{P}_h(\epsilon) \triangleq \{\mathbf{p} \in \mathbb{R}_+^{|\mathcal{J}|} : h(\mathbf{p}) \leq h(\mathbf{p}^*) + \epsilon\}$ for $\epsilon > 0$, where $\mathbf{p}^* \in \mathcal{P}^*$ is some market-clearing price.

Let there be $\epsilon > 0$. By continuity of $h(\mathbf{p})$ and boundedness of $\mathcal{P}_h(\epsilon)$ for any ϵ, there exists $\epsilon_1 > 0$ such that $\mathcal{P}_h(\epsilon_1) \subset B_\epsilon(\mathcal{P})$. Thus, for any $\mathbf{p} \notin B_\epsilon(\mathcal{P})$, we have that $\mathbf{p} \notin \mathcal{P}_h(\epsilon_1)$.

Since $h(\mathbf{p})$ is convex and $\mathcal{P}^*, \mathcal{P}_h(\epsilon)$ are bounded, Proposition 1 in Deng (1998) implies that for $\epsilon_1 > 0$, there exists $\gamma_1 > 0$, such that $\forall \mathbf{p} \notin \mathcal{P}_h(\epsilon_1)$,

$$h(\mathbf{p}) - h(\mathbf{p}^*) - \epsilon_1 \geq \gamma_1 \cdot \mathrm{dist}(\mathbf{p}, \mathcal{P}_h(\epsilon_1)) \,. \tag{34}$$

We will show that the desired inequality (33) holds for $\gamma \triangleq \min\{\gamma_1, \frac{\epsilon_1}{\mathrm{dist}(\mathcal{P}^*, \mathcal{P}_h^c(\epsilon_1))}\} > 0$, where $\mathrm{dist}(\mathcal{P}^*, \mathcal{P}_h^c(\epsilon_1)) \triangleq \inf_{\mathbf{p} \in \mathcal{P}^*, \mathbf{p}' \notin \mathcal{P}_h(\epsilon_1)} ||\mathbf{p} - \mathbf{p}'||_2$.

Note that $\mathrm{dist}(\mathcal{P}^*, \mathcal{P}_h^c(\epsilon)) > 0$ for any $\epsilon > 0$. It follows from the triangle inequality that

$$\mathrm{dist}(\mathbf{p}, \mathcal{P}_h(\epsilon_1)) \geq \mathrm{dist}(\mathbf{p}, \mathcal{P}^*) - \mathrm{dist}(\mathcal{P}^*, \mathcal{P}_h^c(\epsilon_1)) \,.$$

We thus have that for any given $\mathbf{p} \notin B_\epsilon(\mathcal{P})$,

$$\begin{aligned} h(\mathbf{p}) - h(\mathbf{p}^*) &\geq \gamma_1 \cdot \mathrm{dist}\big(\mathbf{p}, \mathcal{P}_h(\epsilon_1)\big) + \epsilon_1 \\ &\geq \gamma\Big(\mathrm{dist}(\mathbf{p}, \mathcal{P}^*) - \mathrm{dist}(\mathcal{P}^*, \mathcal{P}_h^c(\epsilon_1))\Big) + \epsilon_1 \\ &\geq \gamma \cdot \mathrm{dist}(\mathbf{p}, \mathcal{P}^*) - \frac{\epsilon_1}{\mathrm{dist}(\mathcal{P}^*, \mathcal{P}_h^c(\epsilon_1))}\mathrm{dist}(\mathcal{P}^*, \mathcal{P}_h^c(\epsilon_1)) + \epsilon_1 \\ &\geq \gamma \cdot \mathrm{dist}(\mathbf{p}, \mathcal{P}^*) \,. \end{aligned}$$

This completes the proof of step 1.

Step 2. We now bound the Lyapunov drift for the Lyapunov functions

$$\tilde{L}(\mathbf{p}) = \inf_{\mathbf{p}' \in \mathcal{P}^*} \left\{ \frac{1}{2} \sum_{j \in \mathcal{J}} \mu_j (p_j - p_j')^2 \right\} \qquad \text{and} \qquad \tilde{V}(\mathbf{p}) \triangleq \sqrt{\tilde{L}(\mathbf{p})} \,,$$

showing that

$$\mathbb{E}[\tilde{L}(\mathbf{p}_{t+1}) | \mathbf{q}_t] - \tilde{L}(\mathbf{p}_t) \leq -c\,(h(\mathbf{p}_t) - h(\mathbf{p}^*)) + c\frac{2+\lambda}{2(1+\lambda)}\Delta \,. \tag{35}$$

Fix $\mathbf{p}_t$ and let $\mathbf{p}_t^* \triangleq \mathrm{argmin}_{\mathbf{p}' \in \mathcal{P}^*} \frac{1}{2} \sum_{j \in \mathcal{J}} \mu_j (p_{j,t} - p_j')^2$ be the closest market-clearing price to $\mathbf{p}_t$ in the weighted norm (note that the minimum is attainable since $\mathcal{P}^*$ is compact).

As an intermediate step, we first show that

$$\mathbb{E}[\bar{L}_t(\mathbf{p}_{t+1}) | \mathbf{q}_t] - \bar{L}_t(\mathbf{p}_t) \leq -c\,(h(\mathbf{p}_t) - h(\mathbf{p}^*)) + c\frac{2+\lambda}{2(1+\lambda)}\Delta \,, \tag{36}$$

for $\bar{L}_t(\mathbf{p}_t) \triangleq \frac{1}{2} \sum_{j \in \mathcal{J}} \mu_j (p_{j,t} - p_{j,t}^*)^2$.

To see that (36) holds, recall that $p_{j,t+1} = p_{j,t} + \frac{c}{\mu_j}(a_{j,t} - g_{j,t} + l_{j,t})$. We have that

$$\begin{aligned}
&\bar{L}(\mathbf{p}_{t+1}) - \bar{L}(\mathbf{p}_t) \\
&= \frac{1}{2}\sum_{j\in\mathcal{J}} \mu_j \left(\frac{c^2}{\mu_j^2}(a_{j,t} - g_{j,t} + l_{j,t})^2 + \frac{2c}{\mu_j}(p_{j,t} - p^*_{j,t})(a_{j,t} - g_{j,t} + l_{j,t}) \right) \\
&= c\langle \mathbf{p}_t - \mathbf{p}^*_t, \mathbf{a}_t - \mathbf{g}_t \rangle + \frac{1}{2}\sum_{j\in\mathcal{J}} \mu_j \left(\frac{c^2}{\mu_j^2}(a_{j,t} - g_{j,t} + l_{j,t})^2 + \frac{2c}{\mu_j}(p_{j,t} - p^*_{j,t}) l_{j,t} \right) \\
&\overset{(a)}{\leq} c\langle \mathbf{p}_t - \mathbf{p}^*_t, \mathbf{a}_t - \mathbf{g}_t \rangle + \frac{1}{2}\sum_{j\in\mathcal{J}} \mu_j \left(\frac{c^2}{\mu_j^2}(a_{j,t} - g_{j,t} + l_{j,t})^2 + \frac{2c}{\mu_j} p_{j,t} \cdot l_{j,t} \right) \\
&= c\langle \mathbf{p}_t - \mathbf{p}^*_t, \mathbf{a}_t - \mathbf{g}_t \rangle \\
&\quad + \frac{1}{2}\sum_{j\in\mathcal{J}} \mu_j \left(\frac{c^2}{\mu_j^2}\left((a_{j,t} - g_{j,t})^2 + l_{j,t}^2\right) + \frac{2c}{\mu_j}\left(p_{j,t} + \frac{c}{\mu_j}(a_{j,t} - g_{j,t}) \right) l_{j,t} \right),
\end{aligned}$$

where inequality (a) follows because $p^*_{j,t} \geq 0$ and $l_{j,t} \geq 0$. Note that $(p_{j,t} + \frac{c}{\mu_j}(a_{j,t} - g_{j,t}))l_{j,t} \equiv 0$.

Recall equation (12), which establishes that

$$\frac{1}{2}\sum_{j\in\mathcal{J}} \frac{c}{\mu_j}\mathbb{E}[(a_{j,t} - g_{j,t})^2 + l_{j,t}^2 | \mathbf{q}_t] \leq \frac{2+\lambda}{2(1+\lambda)}\Delta .$$

By taking expectations and plugging the above into the RHS, we have

$$\mathbb{E}[\bar{L}(\mathbf{p}_{t+1})|\mathbf{q}_t] - \bar{L}(\mathbf{p}_t) \leq \; c\mathbb{E}[\langle \mathbf{p}_t - \mathbf{p}^*, \mathbf{a}_t - \mathbf{g}_t \rangle | \mathbf{q}_t] + c\frac{2+\lambda}{2(1+\lambda)}\Delta .$$

Using the fact that $\mathbb{E}[\mathbf{a}_t - \mathbf{g}_t | \mathbf{q}_t] \in -\partial h(\mathbf{p}_t)$, we have

$$\mathbb{E}[\bar{L}(\mathbf{p}_{t+1})|\mathbf{q}_t] - \bar{L}(\mathbf{p}_t) \leq \; -c\left(h(\mathbf{p}_t) - h(\mathbf{p}^*)\right) + c\frac{2+\lambda}{2(1+\lambda)}\Delta . \tag{37}$$

To see that the bound (36) for $\bar{L}$ implies the desired bound (35) for $\tilde{L}$, observe that $\bar{L}_t(\mathbf{p}_t) = \tilde{L}(\mathbf{p}_t)$ (by definition), and that $\tilde{L}(\mathbf{p}_{t+1}) \leq \bar{L}_t(\mathbf{p}_{t+1})$.

Step 3. Let $\epsilon > 0$. By the previous steps, there exists γ such that for all $\mathbf{p} \notin B_\epsilon(\mathcal{P}^*)$ we have

$$\begin{aligned}
\mathbb{E}[\tilde{L}(\mathbf{p}_{t+1})|\mathbf{q}_t] - \tilde{L}(\mathbf{p}_t) &\leq \; -c\left(h(\mathbf{p}_t) - h(\mathbf{p}^*)\right) + c\frac{2+\lambda}{2(1+\lambda)}\Delta \\
&\leq \; -c\gamma ||\mathbf{p}_t - \mathbf{p}^*||_2 + c\frac{2+\lambda}{2(1+\lambda)}\Delta .
\end{aligned}$$

Because $f(x) = \sqrt{x}$ is concave for $x \geq 0$, we have that for $y > x > 0$, $f(y) - f(x) \leq (y-x)f'(x) = \frac{y-x}{2\sqrt{x}}$. Therefore,

$$\tilde{V}(\mathbf{p}_{t+1}) - \tilde{V}(\mathbf{p}_t) \leq \frac{\tilde{L}(\mathbf{p}_{t+1}) - \tilde{L}(\mathbf{p}_t)}{2\tilde{V}(\mathbf{p}_t)}\,. \tag{38}$$

Taking conditional expectations given $\mathbf{q}_t$ on both sides of (38) and plugging in (35), we have

$$\begin{aligned}\mathbb{E}[\tilde{V}(\mathbf{p}_{t+1})|\mathbf{q}_t] - \tilde{V}(\mathbf{p}_t) &\leq \frac{\mathbb{E}[\tilde{L}(\mathbf{p}_{t+1})|\mathbf{q}_t] - \tilde{L}(\mathbf{p}_t)}{2\tilde{V}(\mathbf{p}_t)} \\ &\leq \, -\frac{c\gamma}{2} + \frac{2+\lambda}{2(1+\lambda)}\frac{c^2}{2\mu_{\min}\tilde{V}(\mathbf{p}_t)}\,.\end{aligned}$$

For $\mathbf{p}_t$ such that

$$\tilde{V}(\mathbf{p}_t) \geq \frac{c}{\mu_{\min}\gamma}\frac{2+\lambda}{1+\lambda}\,,$$

we have

$$\mathbb{E}[\tilde{V}(\mathbf{p}_{t+1})|\mathbf{q}_t] - \tilde{V}(\mathbf{p}_t) \leq \, -\frac{c\gamma}{4}\,.$$

Now we use a concentration bound from Bertsimas et al. (2001) to prove the desired result. $\tilde{V}(\cdot)$ is a Lyapunov function with exception parameter $\frac{c}{\mu_{\min}\gamma}\frac{2+\lambda}{1+\lambda}$ and negative drift $\frac{c\gamma}{4}$. In each step, the Lyapunov function can increase by at most $\frac{c}{\mu_{\min}}$. Using Theorem 1 in Bertsimas et al. (2001), we have for any $r = 0, 1, \ldots,$

$$\mathbb{P}\left(\tilde{V}\left(\mathbf{p}(c_\ell)\right) > \frac{c}{\mu_{\min}\gamma}\frac{2+\lambda}{1+\lambda} + 2r\frac{c}{\mu_{\min}}\right) \leq \left(\frac{\frac{c}{\mu_{\min}}}{\frac{c}{\mu_{\min}} + \frac{c\gamma}{4}}\right)^{r+1} = \left(\frac{1}{1+\frac{\gamma\mu_{\min}}{4}}\right)^{r+1}.$$

If $\mathbf{p} \notin B_\epsilon(\mathcal{P}^*)$ we have that $\tilde{V}(\mathbf{p}) \geq \frac{\mu_{\min}}{2}\epsilon$. Choosing $c_0 = \mu_{\min}^2\epsilon\gamma/8$ we have that $\Delta_\ell = c_\ell/\mu_{\min} \leq \mu_{\min}\epsilon\gamma\frac{1+\lambda}{4(2+\lambda)}$ for any $c_\ell < c_0$. Plugging in $r = \frac{\mu_{\min}\epsilon}{8\Delta_\ell}$ we have

$$\begin{aligned}\mathbb{P}(\mathbf{p}(c_\ell) \notin B_\epsilon(\mathcal{P})) &\leq \mathbb{P}\left(\tilde{V}(\mathbf{p}(c_\ell)) > \frac{\mu_{\min}}{2}\epsilon\right) \\ &\leq \mathbb{P}\left(\tilde{V}(\mathbf{p}(c_\ell)) > \frac{c}{\mu_{\min}\gamma}\frac{2+\lambda}{1+\lambda} + 2r\frac{c}{\mu_{\min}}\right) \\ &\leq \left(\frac{1}{1+\frac{\gamma\mu_{\min}}{4}}\right)^{\frac{\mu_{\min}\epsilon}{8\Delta_\ell}} \\ &= \exp\left(-\log\left(1+\frac{\gamma\mu_{\min}}{4}\right)\left(\frac{\mu_{\min}\epsilon}{8\Delta_\ell}\right)\right)\,.\end{aligned}$$

This completes the proof.

□

Proof of Corollary 3. Let κ_j denote the probability that an arriving item j is discarded. Denote $\varepsilon = \frac{1}{2}\inf\{p_j \mid \mathbf{p} \in \mathcal{P}^*\}$. We have that if $\mathbf{p} \in B_\varepsilon(\mathcal{P}^*)$, then $p_j \geq \varepsilon > 0$.

For sufficiently small $c_1 > 0$ we have that $\mathbb{E}[c(w_j)|q_j = 0] = c/\mu_j < \varepsilon$ for every $c < c_1$. Thus, $\mathbf{p}_t \in B_\varepsilon(\mathcal{P}^*)$ and $c < c_1$ imply $q_{jt} > 0$; that is, the j-th queue is nonempty in epoch t and an item arriving in epoch t is assigned to some agent. By Proposition 3, there exist $\alpha, \beta, c_0 > 0$, such that for any $c_\ell < \min\{c_0, c_1\}$ we have that

$$\kappa \leq \mathbb{P}\Big(\mathbf{p}(c_\ell) \notin B_\varepsilon(\mathcal{P}^*)\Big) < \beta \cdot e^{-\alpha/\Delta_\ell}\,.$$

□

C Auxiliary Results

This appendix provides auxiliary results to complement our analysis. Appendix C.1 bounds the adjustment size for nonlinear waiting cost functions. Appendix C.2 provides a lemma that tightens the bound of Proposition 3.

C.1 Adjustment size of nonlinear waiting cost functions

We stated our main result (Theorem 1) in terms of the adjustment size Δ. This section shows how Δ is related to the waiting cost function $c(w)$ for nonlinear waiting costs. We focus on the waiting costs that satisfy the assumption below.

Assumption 1. *We consider the following classes of waiting cost functions.*

- Convex waiting costs. *$c(w)$ is convex, twice-differentiable for $w \geq 0$, and $c'(w)$ and $c''(w)$ are subexponential; i.e., there exists α such that $c'(w)$, $c''(w) \leq e^{\alpha w}$ for all $w \geq 0$.*
- Concave waiting costs. *$c(w)$ is concave and twice-differentiable for $w \geq 0$.*

Proposition 5. *Consider the asymptotic regime in Corollary 2 and waiting cost functions satisfying Assumption 1. The following holds:*

1. *For convex $c(w)$, there exists $\ell_0 < \infty$ such that for $\ell \geq \ell_0$, $\Delta \leq \frac{2c'(c^{-1}(v_{\max}))}{\ell\mu_{\min}}$.*
2. *For concave $c(w)$, for any $\ell > 0$, $\Delta \leq \frac{c'(0)}{\ell\mu_{\min}}$.*

Proof of Proposition 5. Consider the system with index ℓ. Let X_t be the interarrival time between the t-th type j item and the $(t+1)$-th type j item. Hence, X_t is an exponential random variable with rate $\ell\mu_j$, and $\{X_t\}_{t=1}^{\infty}$ are i.i.d. Let $S_n \triangleq \sum_{t=1}^{n} X_t$.

Let $q_{\max,\ell}$ be the threshold queue length above which no arriving agent will join that queue. Then approximately

$$v_{\max} = p_j(q_{\max,\ell}) = \mathbb{E}[c(S_{q_{\max,\ell}})] .$$

For convex cost function $c(w)$, by Jensen's inequality we have

$$\mathbb{E}[c(S_{q_{\max,\ell}})] \geq c\left(\mathbb{E}[S_{q_{\max,\ell}}]\right) = c\left(\frac{q_{\max,\ell}}{\ell\mu_j}\right) .$$

Comparing the above two inequalities, we have $q_{\max,\ell} \leq \ell\mu_j c^{-1}(v_{\max})$. Notice that

$$p_j(q_j+1) - p_j(q_j) = \mathbb{E}[c(S_{q_j} + X_{q_j+1}) - c(S_{q_j})] \leq \mathbb{E}[c'(S_{q_j} + X_{q_j+1}) \cdot X_{q_j+1}] ,$$

where the inequality follows from the convexity of $c(w)$. Take the supremum over all $0 \leq q_j \leq q_{\max,\ell}$ on both sides of the above inequality. Because $c(w)$ is convex, $c'(w)$ must be nondecreasing, and hence

$$\Delta = \sup_{0\leq q_j \leq q_{\max,\ell}} \left(p_j(q_j+1) - p_j(q_j)\right) \leq \mathbb{E}[c'(S_{q_{\max,\ell}} + X_{q_{\max,\ell}+1}) \cdot X_{q_{\max,\ell}+1}] .$$

Using Holder's inequality, we have

$$\mathbb{E}[c'(S_{q_{\max,\ell}} + X_{q_{\max,\ell}+1}) \cdot X_{q_{\max,\ell}+1}] \leq \sqrt[\alpha]{\mathbb{E}[(c'(S_{q_{\max,\ell}+1}))^{\alpha}]} \cdot \sqrt[\beta]{\mathbb{E}[X_1^{\beta}]} ,$$

where $\alpha, \beta \in (1,\infty)$ and $\frac{1}{\alpha} + \frac{1}{\beta} = 1$. Because $c(w)$ satisfies Assumption 1, for any $\alpha \in (1,\infty)$ we can apply Lemma 9 and it follows that

$$\lim_{\ell\to\infty} \sqrt[\alpha]{\mathbb{E}[(c'(S_{q_{\max,\ell}}))^{\alpha}]} = \sqrt[\alpha]{c'\left(\mathbb{E}[S_{q_{\max,\ell}}]\right)^{\alpha}} = c'\left(\frac{q_{\max,\ell}}{\ell\mu_j}\right) .$$

Therefore,

$$\lim_{\ell\to\infty} \mathbb{E}[c'(S_{q_{\max,\ell}} + X_{q_{\max,\ell}+1}) \cdot X_{q_{\max,\ell}+1}] \leq c'\left(\frac{q_{\max,\ell}}{\ell\mu_j}\right) \inf_{\beta>1} \sqrt[\beta]{\mathbb{E}[X_1^{\beta}]}$$
$$\leq c'\left(\frac{q_{\max,\ell}}{\ell\mu_j}\right) \mathbb{E}[X_1] = \frac{c'\left(\frac{q_{\max,\ell}}{\ell\mu_j}\right)}{\ell\mu_j} .$$

As a result, there exists $\ell_0 > 0$ such that for $\ell \geq \ell_0$, it holds that

$$\Delta \leq \frac{2c'\left(\frac{q_{\max,\ell}}{\ell\mu_j}\right)}{\ell\mu_j} \leq \frac{2c'\left(c^{-1}(v_{\max})\right)}{\ell\mu_j}\,. \tag{39}$$

By the concavity of $c(w)$, we have that for any $0 \leq q_j \leq q_{\max,\ell}$,

$$p_j(q_j+1) - p_j(q_j) = \mathbb{E}[c(S_{q_j} + X_{q_j+1}) - c(S_{q_j})] \leq \mathbb{E}[c'(0)X_1] = \frac{c'(0)}{\ell\mu_j}\,. \tag{40}$$

Combining (39) and (40) completes the proof. □

Lemma 9. *Let $\{X_i\}_{i=1}^{\infty}$ be i.i.d. exponential random variables with rate 1, $\bar{X}_n \triangleq \frac{1}{n}\sum_{i=1}^{n} X_i$, and $\alpha \in (0,\infty)$. Let $C_1, C_2 \in (0,\infty)$ and let $f(x): \mathbb{R}^+ \to \mathbb{R}$ be a continuously differentiable function such that $f(x) \leq C_1 e^{C_2 x}$, $f'(x) \leq C_1 e^{C_2 x}$ for all $x \in \mathbb{R}^+$. We have*

$$\lim_{N\to\infty} \mathbb{E}\left[f(\alpha \cdot \bar{X}_N)\right] = f(\alpha)\,.$$

Proof of Lemma 9. The result simply follows the proof of Theorem 1(c) in Kozakiewicz et al. (1947), and therefore we omit the details. □

C.2 The rate of convergence of waiting costs

This section presents additional results that provide explicit constants for the rate of convergence shown in Proposition 3. We consider discrete economies with unique market-clearing, and show that prices are very likely to be close to $\mathbf{p}^*$ by showing that prices that deviate from $\mathbf{p}^*$ quickly adjust back. The rate of price adjustment is related to the "sharpness" of the dual objective (5).

The proof of Proposition 3 shows that if for $\gamma(\mathbf{v}, \boldsymbol{\lambda}, \boldsymbol{\mu}) > 0$ it holds that

$$h(\mathbf{p}) - h(\mathbf{p}^*) \geq \gamma(\mathbf{v}, \boldsymbol{\lambda}, \boldsymbol{\mu})||\mathbf{p} - \mathbf{p}^*||_2\,, \tag{41}$$

for any $\mathbf{p} \in \mathbb{R}^{|\mathcal{J}|}$, then the following bound holds:

$$\mathbb{P}(\mathbf{p}(c_\ell) \notin B_\epsilon(\mathcal{P})) \leq \exp\left(-\log\left(1 + \frac{\gamma\mu_{\min}}{4}\right)\left(\frac{\mu_{\min}\epsilon}{8\Delta_\ell}\right)\right)\,.$$

The following lemma provides an explicit lower bound for such γ.

Lemma 10 (Geometry of dual function)**.** *There exists $\gamma(\mathbf{v}, \boldsymbol{\lambda}, \boldsymbol{\mu}) > 0$ such that for any $\mathbf{p} \in \mathbb{R}^{|\mathcal{J}|}$, we have*

$$h(\mathbf{p}) - h(\mathbf{p}^*) \geq \gamma(\mathbf{v}, \boldsymbol{\lambda}, \boldsymbol{\mu})||\mathbf{p} - \mathbf{p}^*||_2\,. \tag{42}$$

Moreover,

$$\gamma(\mathbf{v},\boldsymbol{\lambda},\boldsymbol{\mu}) \geq \frac{1}{\lambda}\left\{\min_{\mathcal{I}\subset\mathcal{J}} \frac{\lambda-\sum_{j\in\mathcal{I}}\mu_j}{\sqrt{|\mathcal{I}|}}\,,\ \min_{\{\mathcal{I}:\mathcal{I}\subset\mathcal{J},\mathcal{I}\supset\cup_{\theta\in\Theta^*}\mathcal{J}^*_\theta\}} \frac{\sum_{j\in\mathcal{I}}\mu_j-\sum_{\theta\in\Theta^*}\lambda_\theta}{\sqrt{|\mathcal{I}|}}\,,\ \min_{j\in\mathcal{J}}\mu_j\right\}.$$

The proof shows that the rate of adjustment is positive for $\mathbf{p}\in\mathcal{P}$ and by the convexity of $h(\mathbf{p})-h(\mathbf{p}^*)$ this holds also for prices not in $\mathcal{P}$.

Proof of Lemma 10. We proceed in two steps.

Step 1. We first show that we can lower bound $h(\mathbf{p})-h(\mathbf{p}^*)$ by a support function:

$$h(\mathbf{p})-h(\mathbf{p}^*) \geq \sup_{s\in\mathcal{S}}\langle\mathbf{p}^*-\mathbf{p},\mathbf{s}\rangle$$

for some convex set $\mathcal{S}$.

For each agent type $\theta\in\Theta$, define

$$\Delta_\theta \triangleq \left\{\mathbf{x}\in\mathbb{R}_+^{|\mathcal{J}|}:\sum_{j\in\mathcal{J}}x_j=1\,,x_j=0\text{ for }j\notin\mathcal{J}^*_\theta\right\},$$

$$\tilde{\Delta}_\theta \triangleq \left\{\mathbf{x}\in\mathbb{R}_+^{|\mathcal{J}|}:\sum_{j\in\mathcal{J}}x_j\leq 1\,,x_j=0\text{ for }j\notin\mathcal{J}^*_\theta\right\},$$

$$\Delta \triangleq \left\{\mathbf{x}\in\mathbb{R}_+^{|\mathcal{J}|}:\sum_{j\in\mathcal{J}}x_j=1\right\},\quad \tilde{\Delta} \triangleq \left\{\mathbf{x}\in\mathbb{R}_+^{|\mathcal{J}|}:\sum_{j\in\mathcal{J}}x_j\leq 1\right\}.$$

Using the definitions above, we can rewrite the dual function (5) as

$$h(\mathbf{p}) = \frac{1}{\lambda}\sum_{\theta\in\Theta}\lambda_\theta\left(\max_{\mathbf{x}_\theta\in\tilde{\Delta}}\sum_{j\in\mathcal{J}}(v(\theta,j)-p_j)x_{\theta,j}\right)+\frac{1}{\lambda}\sum_{j\in\mathcal{J}}\mu_j p_j\,.$$

Let $\mathbf{x}^*_\theta$ be a maximizer of the inner maximization problem described above. Define $\mathbf{s}\in\mathbb{R}^{|\mathcal{J}|}$, where

$$s_j=\sum_{\theta\in\Theta}\lambda_\theta\cdot x^*_{\theta,j}-\mu_j\,;$$

then it is easy to see that $-\frac{1}{\lambda}\mathbf{s}$ is a subgradient of $h(\mathbf{p})$ at $\mathbf{p}$, denoted by $-\frac{1}{\lambda}\mathbf{s}\in\partial h(\mathbf{p})$.

For $\theta \in \Theta^*$, let

$$\mathbf{x}'_\theta \triangleq \text{argmax}_{\mathbf{x}_\theta \in \Delta_\theta} \sum_{j \in \mathcal{J}} (p^*_j - p_j) x_{\theta,j} \,,$$

and for $\theta \in \Theta \backslash \Theta^*$, let

$$\mathbf{x}'_\theta \triangleq \text{argmax}_{\mathbf{x}_\theta \in \tilde{\Delta}_\theta} \sum_{j \in \mathcal{J}} (p^*_j - p_j) x_{\theta,j} \,.$$

The interpretation of $\mathbf{x}'_\theta$ is as follows. Consider an agent of type θ. If the current price is exactly $\mathbf{p}^*$, then the agent is indifferent between the items in $\mathcal{J}^*_\theta$, and strictly prefers these items to other items. If the price deviates a little from $\mathbf{p}^*$: (1) if $\theta \in \Theta^*$, the agent will prefer the item in $\mathcal{J}^*_\theta$ that is the cheapest; (2) if $\theta \notin \Theta^*$, the agent's optimal utility is zero and she will choose an item in $\mathcal{J}^*_\theta$ that is the cheapest only if the price is lower than the optimal price, and she will not choose any item otherwise. Finally, $\mathbf{x}'_\theta$ characterizes the choice of an agent when $\mathbf{p}$ is sufficiently close to $\mathbf{p}^*$.

A key observation is that when $\mathbf{p} \in \mathcal{P}$, the above observation for "sufficiently close" $\mathbf{p}$ holds. Therefore, for $\mathbf{s}' \in \mathbb{R}^{|\mathcal{J}|}$ defined as

$$s'_j \triangleq \sum_{\theta \in \Theta} \lambda_\theta \cdot x'_{\theta,j} - \mu_j \,,$$

we have that $-\frac{1}{\lambda}\mathbf{s}' \in \partial h(\mathbf{p})$ and that, for $\mathbf{p} \in \mathcal{P}^*$,

$$\begin{aligned}
\lambda h(\mathbf{p}) = & \sum_{\theta \in \Theta} \lambda_\theta \sum_{j \in \mathcal{J}} (v(\theta, j) - p_j) x'_{\theta,j} + \sum_{j \in \mathcal{J}} \mu_j p_j \\
= & \sum_{\theta \in \Theta} \lambda_\theta \sum_{j \in \mathcal{J}} (p^*_j - p_j) x'_{\theta,j} + \sum_{j \in \mathcal{J}} \mu_j (p_j - p^*_j) \\
& + \sum_{\theta \in \Theta} \lambda_\theta \sum_{j \in \mathcal{J}} (v(\theta, j) - p^*_j) x'_{\theta,j} + \sum_{j \in \mathcal{J}} \mu_j p^*_j \,.
\end{aligned}$$

Note that the sum of the terms in the last row is exactly $\lambda h(\mathbf{p}^*)$. This is because for an agent of type $\theta \in \Theta^*$, under price $\mathbf{p}$ she must choose an item from $\mathcal{J}^*_\theta$. Hence

$$\sum_{j \in \mathcal{J}} (v(\theta, j) - p^*_j) x'_{\theta,j} = \max_{j \in \mathcal{J}} \, (v(\theta, j) - p^*_j) \,,$$

whereas for an agent of type $\theta \notin \Theta^*$, she either chooses an item from $\mathcal{J}^*_\theta$, or she balks.

Hence

$$\sum_{j\in\mathcal{J}}(v(\theta,j)-p_j^*)x'_{\theta,j}=0=\max_{j\in\mathcal{J}}\,(v(\theta,j)-p_j^*)\,.$$

Thus we have

$$\begin{aligned}\lambda(h(\mathbf{p})-h(\mathbf{p}^*)) = &\sum_{\theta\in\Theta^*}\lambda_\theta\left(\max_{\mathbf{x}_\theta\in\Delta_\theta}\sum_{j\in\mathcal{J}}(p_j^*-p_j)x_{\theta,j}\right)\\ &+\sum_{\theta\in\Theta\backslash\Theta^*}\lambda_\theta\left(\max_{\mathbf{x}_\theta\in\tilde{\Delta}_\theta}\sum_{j\in\mathcal{J}}(p_j^*-p_j)x_{\theta,j}\right)+\sum_{j\in\mathcal{J}}\mu_j(p_j-p_j^*)\,. \end{aligned} \tag{43}$$

Define the *rate region* $\mathcal{S}$ as:

$$\mathcal{S}\triangleq\left\{\sum_{\theta\in\Theta}\lambda_\theta\mathbf{x}_\theta-\boldsymbol{\mu}:\mathbf{x}_\theta\in\Delta_\theta\text{ for }\theta\in\Theta^*\,,\mathbf{x}_\theta\in\tilde{\Delta}_\theta\text{ for }\theta\in\Theta\backslash\Theta^*\right\}\,,$$

which is the set of possible rates of change of dual prices when $\mathbf{p}\in\mathcal{P}$. Therefore we can rewrite the RHS of (43) as

$$\sup_{\mathbf{s}\in\mathcal{S}}\,\langle\mathbf{p}^*-\mathbf{p},\mathbf{s}\rangle\,.$$

Using the fact that $h(\mathbf{p})$ is convex, we have for any $\mathbf{p}$,

$$h(\mathbf{p})-h(\mathbf{p}^*)\geq\frac{1}{\lambda}\sup_{\mathbf{s}\in\mathcal{S}}\,\langle\mathbf{p}^*-\mathbf{p},\mathbf{s}\rangle\,.$$

This completes step 1.

Step 2. Characterizing the set $\mathcal{S}$. Note that $\mathcal{S}$ is the Minkowski sum of simplices shifted by $\boldsymbol{\mu}$, which is known as the generalized permutohedron (see, e.g., Postnikov, 2009). Using Proposition 6.3 from Postnikov (2009), we have the following defining inequalities of $\mathcal{S}$:

$$\begin{aligned} \sum_{j\in\mathcal{I}}s_j &\leq\lambda-\sum_{j\in\mathcal{I}}\mu_j\,, && \forall\mathcal{I}\subset\mathcal{J}\,,\\ \sum_{j\in\mathcal{I}}s_j &\geq\sum_{\theta\in\Theta^*}\lambda_\theta-\sum_{j\in\mathcal{I}}\mu_j\,, && \forall\mathcal{I}:\mathcal{I}\subset\mathcal{J}\,,\mathcal{I}\supset\cup_{\theta\in\Theta^*}\mathcal{J}_\theta^*\,,\\ s_j &\geq\,-\mu_j\,, && \forall j\in\mathcal{J}\,. \end{aligned}$$

We first argue that there exists $\epsilon>0$ such that the ball $B(\mathbf{0},\epsilon)$ is contained in $\mathcal{S}$. This can be proved by contradiction: if it is not true, then using (43), we can show that the minimizer of $h(\mathbf{p})$ is nonunique, leading to a contradiction of the assumption that the dual optimum is unique. Note that this already leads to a lower bound of $h(\mathbf{p}) - h(\mathbf{p}^*)$: we have $\epsilon\frac{\mathbf{p}^*-\mathbf{p}}{||\mathbf{p}^*-\mathbf{p}||_2} \subset \mathcal{S}$, and hence $\lambda(h(\mathbf{p}) - h(\mathbf{p}^*)) \geq \epsilon||\mathbf{p}^* - \mathbf{p}||_2$.

It remains to quantitatively characterize ϵ. To simplify the notation, we consider the centered version of $\mathbf{p}$, defined as $\tilde{\mathbf{p}} \triangleq \mathbf{p}^* - \mathbf{p}$; let $\tilde{h}(\tilde{\mathbf{p}}) \triangleq h(\mathbf{p}) - h(\mathbf{p}^*)$.

Since $\mathcal{S}$ is defined "locally" (i.e., for $\mathbf{p} \in \mathcal{P}$), all the arguments below assume that $\mathbf{p} \in \mathcal{P}$. We have derived that $\tilde{h}(\tilde{\mathbf{p}}) = \sup_{\mathbf{s}\in\mathcal{S}} \langle \tilde{\mathbf{p}}, \mathbf{s}\rangle$. Define the level sets of $\tilde{h}(\tilde{\mathbf{p}})$:

$$\mathcal{L} \triangleq \left\{\tilde{\mathbf{p}} \in \mathbb{R}^{|\mathcal{J}|} : \tilde{p}_j \leq 0 \text{ for } j \neq \mathcal{J}^*, \tilde{h}(\tilde{\mathbf{p}}) \leq 1\right\} .$$

Here the constraints $\tilde{p}_j \leq 0$ for $j \neq \mathcal{J}^*$ come from the fact that $\mathbf{p} \geq \mathbf{0}$. Using the theory of polar duality (see, e.g., Gallier, 2008), since the ball $B(\mathbf{0}, \epsilon)$ is contained in $\mathcal{S}$, we have

$$(\mathcal{S} \cap \{\tilde{\mathbf{p}} : \tilde{p}_j \leq 0 \text{ for } j \neq \mathcal{J}^*\})^* \subset (B(\mathbf{0}, \epsilon) \cap \{\tilde{\mathbf{p}} : \tilde{p}_j \leq 0 \text{ for } j \neq \mathcal{J}^*\})^* .$$

Here the asterisks outside the parentheses stand for polar sets. Denote the LHS set as $\mathcal{B}^*$ and the RHS set as $\mathcal{L}^*$. We have

$$\mathcal{B}^* = \left\{\mathbf{s} + \sum_{j\notin\mathcal{J}^*} \gamma_j \mathbf{e}_j : ||\mathbf{s}||_2 \leq \epsilon\, , \gamma_j \geq 0\right\} ,$$
$$\mathcal{L}^* = \left\{\mathbf{s} + \sum_{j\notin\mathcal{J}^*} \gamma_j \mathbf{e}_j : \mathbf{s} \in \mathcal{S}\, , \gamma_j \geq 0\right\} .$$

Because $\mathcal{B}^* \subset \mathcal{L}^*$, ϵ can take a value up to the inradius of $\mathcal{S}$, which is larger than the minimum of the distances between $\mathbf{0}$ and the defining hyperplanes of $\mathcal{S}$. It follows that

$$\epsilon \geq \left\{\min_{\mathcal{I}\subset\mathcal{J}} \frac{\lambda - \sum_{j\in\mathcal{I}} \mu_j}{\sqrt{|\mathcal{I}|}}\, , \min_{\{\mathcal{I}:\mathcal{I}\subset\mathcal{J}\, ,\mathcal{I}\supset\cup_{\theta\in\Theta^*}\mathcal{J}^*_\theta\}} \frac{\sum_{j\in\mathcal{I}} \mu_j - \sum_{\theta\in\Theta^*} \lambda_\theta}{\sqrt{|\mathcal{I}|}}\, ,\ \min_{j\in\mathcal{J}} \mu_j\right\} .$$

□