Price Discovery in Waiting Lists: A Connection to Stochastic Gradient Descent

Preliminary Version - Please Do Not Circulate*

Itai Ashlagi Jacob D. Leshno Pengyu Qian Amin Saberi

Abstract

Waiting lists allocate items by offering agents a choice among items with associated waiting times. These waiting times serve as prices that are formed by the stochastic arrivals and departures of agents. We study the allocative efficiency under such dynamically adjusting prices by drawing a connection between this price adjustment process and the stochastic gradient descent optimization algorithm. We show that the loss due to price fluctuations is bounded by the granularity of price changes. Additional conditions allow us to identify markets where the loss is close to the bound or exponentially small. Our results show that a simple price adjustment heuristic can perform well, but may be slow to adjust to changes in demand because of a trade-off between the speed of adaptation and loss from price fluctuations.

1 Introduction

Public scarce resources are often allocated through waiting lists, in which agents can select which resource to wait for For example, the New York Public Housing Authority asks applicants to choose a project specific queue. If utilities are quasi-linear in waiting costs and total waiting costs are constant across assignments, waiting times act as prices that clear the market and create an efficient assignment. In this sense, waiting lists are similar to standard competitive equilibrium (CE) models, with the exception that prices are quoted in waiting times instead of monetary transfers. Waiting lists mechanisms have a natural price formation process. Indeed, waiting times, or prices, are naturally adjusted with arrival of agents and resources. This paper is concerned with how the price dynamics that are inherent to waiting lists impact allocative efficiency.

*A one page abstract of this work appeared in the Proceedings of the 21st ACM Conference on Economics and Computation under the name "Queue lengths as constantly adapting prices: Allocative efficiency under random dynamics"

The economy considered has items of different types as well as agents arriving over time. A waiting list mechanism maintains one observable queue for each item type. Agents have unit demand, heterogeneous private values for items. The utility of an agent is quasi-linear in waiting costs. Upon arrival, the agent, who maximizes expected utility, chooses either a single queue and waits until she is assigned an item, or, leaves immediately and receives an outside option. When an item arrives, it is assigned to an agent in the respective queue or discarded if that queue is empty.

Prices in this economy depend on the current state of the queues; the price of an item type decreases when such an item is assigned to an agent, and increases when an agent joins the corresponding queue. This means that different agents may face different prices upon arrival. This stochastic evolution creates a challenge to analyze this simple economy.

To understand how price adjustments affect allocative efficiency, it is helpful to first relate the economy to a canonical CE model. Suppose that instead of dynamically assigning items, a center could wait for a large time frame and then simultaneously assign all accumulated items to the agents By the first welfare theorem, prices that clear the market and form a CE will generate an assignment of items to agents that maximizes the total value agents ascribe to their assigned items. Assuming that total waiting costs are constant across assignments, the value generated by the CE assignment gives an upper bound on the welfare in the dynamic model. Moreover, this welfare can be achieved if prices (expected waiting costs at each queue) remain constant at CE prices.

Since prices in the economy fluctuate, they may fail to generate the same allocative efficiency as CE prices. To illustrate this consider a simple example with a single type of item. Agents' valuations are distributed uniformly between 0 and 1 and they arrive twice as fast as items. Since only half the agents can be assigned, it is optimal to set a price equal to $1/2$ and assign the item to the agents with valuation above that price. But in the waiting list the price fluctuates below and above $1/2$. This results in some agents with lower values being assigned and some agents with high values selecting the outside option.

With multiple types of items, prices follow a multidimensional Markov process. While this process is easy to describe, the stationary distribution cannot be derived tractably. For this reason, several related papers restricted attention to economies with at most two item types (see, e.g., Baccara et al. (2018); Leshno (2017)). However, we can bound the welfare loss in the economy by drawing a connection between the price dynamics under the waiting list mechanism, and the stochastic gradient descent (SGD) optimization algorithm. The main insight is that while price adjustments are stochastic, prices tend to adjust towards their CE levels. Formally, the expected price adjustment is a subgradient of a dual of the welfare maximization problem.

A key quantity is the granularity of price adjustments Δ, defined as the maximal increase

in price (waiting costs) due to the addition of one agent to a queue. Equivalently, Δ determines the size of the adjustment step taken by the SGD algorithm after seeing one sample. This granularity is a key distinction between the price adjustment process in our economy and standard usage of SGD in optimization: to make the SGD converge the step size is reduced to zero over time, while in waiting lists the step size is exogenously given by agent sensitivity to waiting costs and remains constant. Despite this distinction, the connection to SGD allows us to leverage tools from Lyapunov theory to obtain tractable bounds for welfare.

The main finding is that the welfare loss in waiting lists is bounded by the granularity of price adjustments Δ times a factor that depends only on the relative arrival rates of agents and items. The dynamic price adjustments keep prices "close" to the CE prices, despite never converging. The welfare loss due to constant fluctuation of prices depends on the magnitude of these fluctuations, which become small as Δ becomes small.

For further intuition, consider again the simple example and assume waiting costs are linear. If waiting behind 10 agents in the queue gives a waiting cost of $1/2$, then each arrival or departure of an agent changes the price by 10% of the item's maximal value. If agents are more patient and waiting behind 100 agents in the queue gives a waiting cost of $1/2$, then each arrival or departure of an agent changes the price by only 1% of the item's maximal value. So Δ is smaller when agents are more patient. Intuitively, prices will fluctuate less when agents are patient resulting in less welfare loss.

Two additional results complement the main finding. We show that the bound is essentially tight (up to constant factors) by explicitly constructing an economy with large allocative efficiency loss. A distinct and important feature of this economy is the multiplicity of CE prices.

Finally, it is shown that the loss is generically much smaller for economies with finitely many agent types. When the number of agent types is finite, the CE prices are unique for all but a zero measure of arrival rates. And when CE prices are unique the welfare loss becomes exponentially small as Δ tends to zero.

The analysis can offer insight for general price adaptation algorithms. Consider a planner who is able to choose the granularity of price adjustments Δ. Increasing Δ will make the prices adjust more quickly if the distribution of agent valuations changes. But a higher value of Δ will also cause losses when agent valuations remains the same, but prices change due to the random arrivals. This tension is inherent to any price adaptation mechanism that can only observe noisy signals, and needs to trade-off overreactions due to imprecise information and slow reaction time due to the need to accumulate sufficiently precise information. A particular feature may be of interest: firms that adjust prices according to an SGD heuristic may find it optimal to react slowly to changes in monetary policy if these changes have

unclear implications for their demand new prices new to be learned through the dynamics.

1.1 Related Literature

The first area this paper is related is dynamic matching motivated by applications in public housing Kaplan (1984, 1988) and organ allocation Zenios (1999). Several papers compare the efficiency across queuing and lottery mechanisms for restricted set of preferences (Su and Zenios, 2004; Bloch and Cantala, 2017; Arnosti and Shi, 2020). Some papers consider optimal design of dynamic allocation without accounting for the stochastic process (Polterovich, 1993; Braverman et al., 2016; Afeche et al., 2019; Ata et al., 2021). Our paper focuses on understanding the relation between the dynamics in the queueing mechanism the economic efficiency.

Several papers are concerned with the stochasticty 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.

This paper also relates to papers that are concerned with 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). These papers consider static market and study identify an adjustment process that converges to equilibrium. In contrast, in the dynamic market considered here "prices" never converge.[1]

Similar ideas have been considered in the network control literature (see, e.g., Georgiadis et al., 2006), but the model consider a much more general settings (continuum of types, nonlinear waiting costs) and establish novel results (exponential loss).

2 A Simple Illustrative Example

We consider a simple market to illustrate the loss from price-fluctuations. There is a single kind of item, which arrives according to a Poisson process with rate 1. Agents arrive according to a Poisson process with rate 2. Arriving agents choose whether to join the queue and wait to receive an item, or be immediately assigned an outside option and receive 0 utility. The utility of an agent who joins the queue is

$$v - c \cdot w$$

[1]Some papers identify price adjustments processes that converge to market clearing prices in a static version of our assignment problem though with finite many types of agents (e.g., Bertsekas (1979); Demange et al. (1986)).

where w is the amount of time the agent waits, $c = 0.02$ is equal to the expected waiting costs of waiting for one item's arrival, and v is the agent's value of receiving the item. Each agent's value v is independently drawn from the uniform distribution on $[0, 1]$. Agents observe the current number of agents in the queue when they arrive. Agents in the queue are assigned in the order in which they joined the queue (First-Come First-Served). When an agent with value v arrive, she will join the queue if $v - c \cdot \mathbb{E}[w|q] > 0$ where $\mathbb{E}[w|q] = q + 1$ is the expected wait given there are currently q agents in the queue.

As a benchmark, consider a planner who collects all agents and items that arrive up to time T and assigns all items at time T. If T is large, the planner collects approximately twice as many agents as items, and the distribution of agent values is approximately $U[0, 1]$. Because only half the agents can be assigned an item, allocative efficiency is maximized by allocating items only to agents with a value $v \in [1/2, 1]$.

Figure 2 shows the distribution of queue length an arriving agent faces. On average, the queue length implies an expected waiting cost equal to $1/2$. If all agents faced a price (in waiting cost) equal to $1/2$, only agents with a value $v \in [1/2, 1]$ would be assigned items and the maximal allocative efficiency would be generated.

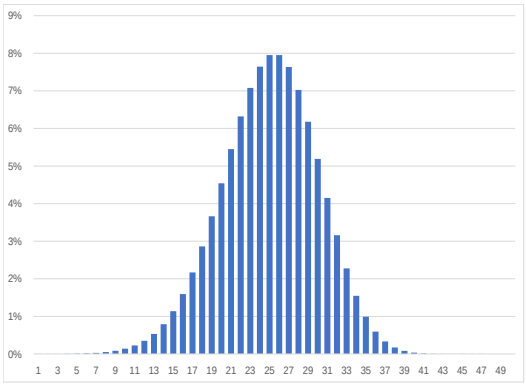


Figure 1: Distribution of the queue length when there is a single item for which agent values are distributed $v \sim U[0, 1]$ and $c = 0.02$.

However, Figure 2 shows there is considerable variation in the prices agents face. As a result, the allocation may be inefficient. An agent with a value of $v = 0.41$ may be assigned if she is lucky to arrive when the queue happens to be short. An agent with a value of $v = 0.59$ may not be assigned if she is arrives when the queue happens to be long. Figure 2 shows the implies probability that an agent is assigned as a function of the agent's value. It shows that agents with a value of $0.41, 0.59$ are assigned the item with probabilities 18% and 81%, respectively.

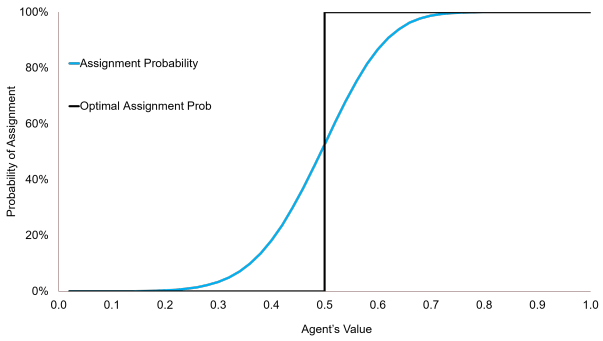


Figure 2: Implied probability of assignment given the agent's value under the optimal assignment and the waiting list mechanism

3 Model

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

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 distribution F over the set of types Θ. We assume that Θ is a compact subset of an Euclidean space, and allow for both finitely many agent types as well as a continuum of agents. We say that there are finitely many agent types if F corresponds to finitely many atoms.

Items arrive according to a Poisson process with total rate 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 $\sum_{j\in\mathcal{J}} \mu_j = 1$. Denote by $\mu_{\min} \triangleq \min_{j\in\mathcal{J}} \mu_j > 0$, $\mu_{\max} \triangleq \max_{j\in\mathcal{J}} \mu_j > 0$. 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 from getting assigned to 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 to type j after waiting w units of time receives a utility of

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

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

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,[2] and $c(0) = 0$, $\lim_{w\to\infty} c(w) = \infty$.

To simplify notation, we consider an equivalent discrete time process[3] indexed by t which 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 arrives. 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 item arriving at t.

Assignments and Allocative Efficiency An *allocation* η assigns each agent with one item, and each non-auxiliary item is assigned to at most one agent. The allocative efficiency of a matching is defined as the average item's value to its assigned agent. Formally, given allocation 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 $A_T = \sum_{t\le T} \xi_t$ be the total number of agents that arrived up to epoch T. Allocative efficiency under η is defined as

$$W(\eta) = \liminf_{T\to\infty} \frac{1}{A_T} \sum_{t=1}^{T} \xi_t v(\theta_t, \eta_t)\,. \tag{1}$$

We restrict attention to allocations that satisfy a *no-Ponzi* condition. Loosely speaking, this condition ensures the assignment is approximately valid if the market terminates at some large finite time.[4] Formally, let $\mathcal{R}_T(\eta)$ denote the number of agents and items that arrived by time T and are waiting to be assigned at time T.[5] 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.

Define optimal allocative efficiency to be

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

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

Allocation by a Waiting List Our main interest is to analyze the allocative efficiency of the allocation generated by a standard waiting list mechanism. The mechanism holds

[2] 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 $c'(w), c''(w) \le e^{\alpha w}$ for all $w \ge 0$.

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

[4] For example, an allocation that assigns all agents to items in the market described in Section 2 is not valid if the market terminates at any finite time, as only (approximately) half the agents can be assigned items.

[5] In other words, $\mathcal{R}_T(\eta)$ counts the number of agents who arrived before time T and are assigned under η to 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.

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 take the auxiliary item immediately (i.e., 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 mechanism, 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 will choose to join the queue for item $a(\theta, \mathbf{q}) \in$ defined by[6]

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

where we define $p_j(q_j) \triangleq \mathbb{E}[c(w_j)|q_j]$. We allow agents to leave without joining any queue, and simplify notation by setting $p_\emptyset(\cdot) \equiv 0$. For notation simplicity, 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 $\mathbf{q}_t$.

Given a realization, let η^{WL} denote the allocation induced by the waiting list. Under our assumptions, η^{WL} satisfies the no-Ponzi condition.[7] We denote the expected allocative efficiency of the waiting list mechanism by

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

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

4 Bounding the Allocative Efficiency Loss

In the waiting list mechanism, each agent is presented a menu of items and associated expected waiting costs. We consider the waiting cost of an item as the item's price, and henceforth refer to the expected waiting cost of item j as the price of item j.

Standard intuition from competitive equilibria suggests that appropriate prices and result in an allocation that maximizes allocative efficiency. But in the waiting list there is no planner that sets prices. Instead, prices are determined by the current state of the queues.

[6] To simplify notation, we assume that the argmax is unique and implicitly rely on a tie-breaking rule to ensure a unique selection if the agent is indifferent between multiple items. Our results do not depend on the choice of tie-breaking rule. For example, agents may randomly choose an item in the argmax or choose the lowest index item within the argmax.

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

Prices adapt over time in a process that is similar to a tâtonnement process: the price 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.

Because prices are state-dependent, prices fluctuate over time. For example, if an agent θ arrives immediately after several copies of the item j happened to arrive, agent θ will face a lower price p_j (i.e., expected waiting cost for item j). The example in Section 2 shows that such stochastic price fluctuations can lead to lower allocative efficiency.

We seek to understand the allocative efficiency of the fluctuating state-dependent prices in the waiting list. A natural approach would be to calculate the stationary distribution of prices. Unfortunately, this stationary distribution is not tractable when $|\mathcal{J}| > 2$, that is, there are strictly more than 2 kinds of items.[8] 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 nonlinear waiting costs.

Our analysis shows that the following attribute plays an central role in determining allocative efficiency:

Definition 1. *The step size Δ is the maximal change in price due to a single arrive and is given by*[9]

$$\Delta \triangleq \max_{j\in\mathcal{J}} \max_{0\le q\le q_{\max}} \big(p_j(q) - p_j(q-1)\big).$$

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, i.e., $c(t) = c \cdot t$ for some $c > 0$, we have $\Delta = c/\mu_{\min}$. That is, Δ is the expected cost of waiting for a single arrival of the least frequent item.

We can now state our main result.

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

$$W^{\mathrm{WL}} \ge W^{\mathrm{OPT}} - \frac{\lambda+2}{2\lambda}\Delta\,. \tag{3}$$

[PQ: I redid the calculation for the case $p = c(q+1)/\mu$ (previously it was $p = cq/\mu$), and the numerator of the coefficient of Δ goes from $\lambda+1$ to $\lambda+2$.]

Theorem 1 states that the queueing mechanism achieves allocative efficiency that is close to optimal in general dynamic markets. As an example, suppose waiting costs are linear and the arrival rate of agents equals the arrival rate of items ($\lambda = \mu = 1$). Then, the

[8] Because of this limitation, previous papers that relied on calculation of the exact stationary distribution of the underlying Markov chain were limited to a model with 2 items (e.g., Leshno (2017); Bloch and Cantala (2017); Baccara et al. (2018)).

[9] Let $p_j(-1) = 0$ so that Δ is well-defined. Note that $p_j(0) > 0$ because even if the agent arrives to an empty queue, she still needs to wait for the next to item to arrive.

allocative efficiency loss is bounded by the cost of waiting for a single arrival of the least frequent item times $\frac{3}{2}$. In Section 4.2 we provide an example giving a lower bound on the allocative efficiency loss.

To gain intuition for Theorem 1, we draw connections to two related problems.

Problem I: Duality for Static Allocation. Consider the static allocation problem in which a planner chooses a static assignment of the expected amount of agents that arrive per unit time to the expected amount of items that arrive per unit time in order to maximize allocative efficiency. We refer to the maximal value per agent the assignment can generate as the optimal static allocative efficiency and denote it by W^*. The value of W^* is the optimal value of assignment problem (4).

$$\begin{aligned} W^* = \max_{\{x_{\theta j}\}_{\theta\in\Theta, j\in\mathcal{J}}} \quad & \sum_{j\in\mathcal{J}} \int_\Theta x_{\theta j}\, v(\theta, j)\, dF(\theta) & & \\ \text{subject to} \quad & \sum_{j\in\mathcal{J}} x_{\theta j} \leq 1,\ x_{\theta j} \in [0,1] & \forall \theta \in \Theta & \qquad (4) \\ & \int_\Theta \lambda x_{\theta j}\, dF(\theta) \leq \mu_j & \forall j \in \mathcal{J} & \end{aligned}$$

In problem (4), $x_{\theta j}$ is the share of agents of type θ that are assigned item j. The first constraint requires that the shares $x_{\theta j}$ are well defined. The second constraint is the resource constraint, it requires that the expected amount of item j arriving per unit time should at least as large as the expected amount of agents that arrive per unit time and are assigned to item j.

Proposition 1. *The optimal allocative efficiency is $W^{\text{OPT}} = W^*$, where W^* is the optimal static allocative efficiency.*

The proof of Proposition 1 is in Appendix A.

It will be useful to consider the dual problem of the assignment problem (4), which optimizes over possible prices. The following strong duality result is well-known and we therefore omit the proof.[10]

Lemma 1 (Monge-Kantorovich duality)**.** *The optimal value W^* of assignment problem (4), coincides with the optimal value of the following dual optimization problem:*

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

[10]Problem (4) is known as the (unbalanced) optimal transport problem, which has the strong duality property stated in Lemma 1. For further details, see, e.g., Galichon (2018).

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

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

We use $\mathbf{p}^*$ to denote some optimal value for the dual problem, and refer to $\mathbf{p}^*$ as optimal prices.

Problem II: Stochastic Gradient Descent. We relate the price adaptation in the waiting list to the run of the stochastic gradient descent optimization algorithm. The SGD algorithm can be regarded as a stochastic version of the gradient descent optimization algorithm. Each step of the SGD is random, but the expected step of SGD corresponds to a step of gradient descent. SGD optimization is commonly used in machine learning, e.g., for training a neural networks (BackProp et al., 1998). By understanding the connection between our problem and the SGD algorithm we are able to leverage the substantial theory on SGD algorithms.

The following lemma establishes the connection between the waiting list and the SGD algorithm:

Lemma 2. *If the system is in state $\mathbf{q}_t$ 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 $\mathbf{p}_t = \mathbf{p}(\mathbf{q}_t)$.*

Proof. The expected adjustment to the length of queue j from a single arrival is

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

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

In other words, Lemma 2 says that the stochastic price adjustment from one arrival corresponds to a step of an SGD algorithm for the dual objective $h(\mathbf{p})$ (Lemma 1). Loosely speaking, the waiting lists adjusts prices in the right direction on average, but the adjustment is random because it depends on the realization of a single arrival. However, in order for gradient descent and SGD algorithms to converge, the step size must decrease to zero as the algorithm approaches the optimal value. In 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.

Intuition for Theorem 1. The main part of the proof uses a Lyapunov potential function to decompose the expected value from the next arrival. Each agent's arrival generates a value of assignment, and the value generated depends on the current state of the queues. This value is related to the dual objective (5) evaluated at the current prices $\mathbf{p}_t$, which by Lemma 1 is at least as high as W^*. Each arrival also changes the current state. We use a Lyapunov function to capture the "potential" given the current price $L(\mathbf{p}_t)$. We decompose the expected value from the next arrival into a combination of the objective, change in potential, and a per-period loss.

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

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

Proof of Theorem 1. The proof adopts the Lyapunov analysis approach. We give the main arguments of the proof for the special case of linear waiting cost here and relegate technical lemmas and the proof of general cases to the appendix.

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

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

We have that

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

We introduce several definitions for the analysis. As is standard in Lyapunov analysis, let the Lyapunov function be the following quadratic function:

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

Let $\mathbf{a}_t$ and $\mathbf{d}_t$ be the vectors representing the arriving agent and item at time t, respectively:

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

Let $u_{j,t} \triangleq \max\{0, d_{j,t} - q_{j,t} - a_{j,t}\}$ denote the number of discarded items of type j at time t.[11] The evolution of the length of queue j is governed by

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

By Lemma 3 we have that

$$\begin{aligned}\mathbb{E}[v(\theta_t, a(\theta_t,\mathbf{q}_t))\xi_t|\mathbf{q}_t] = \frac{\lambda}{1+\lambda}h(\mathbf{p}_t) - \frac{1}{c}\left(L(\mathbf{p}_t) - \mathbb{E}[L(\mathbf{p}_{t+1}) \mid \mathbf{q}_t]\right) \\ - \frac{1}{2}\sum_{j\in\mathcal{J}}\frac{c}{\mu_j}\mathbb{E}[(a_{j,t}-d_{j,t})^2 + u_{j,t}^2 \mid \mathbf{q}_t]\end{aligned} \tag{8}$$

By Lemma 4 we have that $\frac{1}{2}\sum_{j\in\mathcal{J}}\frac{c}{\mu_j}\mathbb{E}[(a_{j,t}-d_{j,t})^2 + u_{j,t}^2 \mid \mathbf{q}_t] \le \frac{2+\lambda}{2(1+\lambda)}\Delta$. By Lemma 1 we have that $h(\mathbf{p}_t) \ge W^*$. Together, we have that

$$\mathbb{E}[v(\theta_t, a(\theta_t,\mathbf{q}_t))\xi_t|\mathbf{q}_t] \ge \ \frac{\lambda}{1+\lambda}W^* - \frac{1}{c}\left(L(\mathbf{p}_t) - \mathbb{E}[L(\mathbf{p}_{t+1}) \mid \mathbf{q}_t]\right) - \frac{2+\lambda}{2(1+\lambda)}\Delta\,.$$

Therefore, we have that

$$\begin{aligned}\mathbb{E}\left[W_T(\eta^{\mathrm{WL}})\right] &= \mathbb{E}\left[\sum_{t=1}^{T}\xi_t\cdot v(\theta_t, a(\theta_t,\mathbf{q}_t))\right] \\ &= \mathbb{E}\left[\mathbb{E}\left[\sum_{t=1}^{T}\xi_t\cdot v(\theta_t, a(\theta_t,\mathbf{q}_t)) \mid \mathbf{q}_t\right]\right] \\ &= \mathbb{E}\left[\sum_{t=1}^{T}\mathbb{E}\left[\xi_t\cdot v(\theta_t, a(\theta_t,\mathbf{q}_t)) \mid \mathbf{q}_t\right]\right] \\ &\ge \mathbb{E}\left[\sum_{t=1}^{T}\frac{\lambda}{1+\lambda}W^* - \frac{1}{c}\left(L(\mathbf{p}_t) - \mathbb{E}[L(\mathbf{p}_{t+1} \mid \mathbf{q}_t]\right) - \frac{2+\lambda}{2(1+\lambda)}\Delta\right] \\ &= T\frac{\lambda}{1+\lambda}W^* - \frac{1}{c}\left(L(\mathbf{p}_1) - \mathbb{E}[L(\mathbf{p}_{T+1})]\right) - T\frac{2+\lambda}{2(1+\lambda)}\Delta\,.\end{aligned} \tag{9}$$

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

By Lemma 5 we can translate the bound for $\mathbb{E}[W_T(\eta^{\mathrm{WL}})]$ to a bound for W^{WL}, i.e.,

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

Plugging in (9) to the above equality, we have

$$W^{\mathrm{WL}} \geq W^* - \frac{2+\lambda}{2\lambda}\Delta .$$

This concludes the proof. □

Lemma 3. *If $c(w) = c \cdot w$, we have that*

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

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

$$\begin{aligned}& L(\mathbf{p}_t) - L(\mathbf{p}_{t+1}) \\ &= \frac{1}{2}\sum_{j\in\mathcal{J}} \mu_j \left[p_{j,t}^2 - \left(p_{j,t} + \frac{c}{\mu_j}(a_{j,t} - d_{j,t} + u_{j,t})\right)^2\right] \\ &= \frac{1}{2}\sum_{j\in\mathcal{J}} \mu_j \left[p_{j,t}^2 - \left(p_{j,t} + \frac{c}{\mu_j}(a_{j,t} - d_{j,t})\right)^2 - \frac{c^2}{\mu_j^2}u_{j,t}^2 - \frac{2c}{\mu_j}\left(p_{j,t} + \frac{c}{\mu_j}(a_{j,t} - d_{j,t})\right) u_{j,t}\right] \\ &\overset{(a)}{=} \frac{1}{2}\sum_{j\in\mathcal{J}} \mu_j \left[p_{j,t}^2 - \left(p_{j,t} + \frac{c}{\mu_j}(a_{j,t} - d_{j,t})\right)^2 - \frac{c^2}{\mu_j^2}u_{j,t}^2\right] ,\end{aligned}$$

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

We further simplify the Lyapunov drift as follows.

$$\begin{aligned}L(\mathbf{p}_t) - L(\mathbf{p}_{t+1}) &= \frac{1}{2}\sum_{j\in\mathcal{J}} \mu_j \left[p_{j,t}^2 - \left(p_{j,t} + \frac{c}{\mu_j}(a_{j,t} - d_{j,t})\right)^2\right] - \frac{1}{2}\sum_{j\in\mathcal{J}} \frac{c^2}{\mu_j} u_{j,t}^2 \\ &= -\frac{1}{2}\sum_{j\in\mathcal{J}} \mu_j \left[\frac{2c}{\mu_j}(a_{j,t} - d_{j,t})p_{j,t} + \frac{c^2}{\mu_j^2}(a_{j,t} - d_{j,t})^2\right] - \frac{1}{2}\sum_{j\in\mathcal{J}} \frac{c^2}{\mu_j} u_{j,t}^2 \\ &= -c\sum_{j\in\mathcal{J}} (a_{j,t} - d_{j,t})p_{j,t} - \frac{1}{2}\sum_{j\in\mathcal{J}} \frac{c^2}{\mu_j}\left((a_{j,t} - d_{j,t})^2 + u_{j,t}^2\right) . \qquad (10)\end{aligned}$$

We expanded the expected value of the next arrival $\mathbb{E}\Big[v(\theta_t, a(\theta_t, \mathbf{q}_t))\xi_t \mid \mathbf{q}_t\Big]$ plus the

expected value of the term $\langle \mathbf{p}_t, \mathbf{a}_t - \mathbf{d}_t \rangle$ to show these are related to the dual objective given in (5).

$$
\begin{aligned}
&\mathbb{E}\Big[v(\theta_t, a(\theta_t, \mathbf{q}_t))\xi_t - \langle \mathbf{p}_t, \mathbf{a}_t - \mathbf{d}_t \rangle \mid \mathbf{q}_t\Big] \\
&= \mathbb{E}\Big[v(\theta_t, a(\theta_t, \mathbf{q}_t))\xi_t - \sum_{j\in\mathcal{J}} p_{j,t}(a_{j,t} - d_{j,t}) \mid \mathbf{q}_t\Big] \\
&= \mathbb{E}\Big[\max_{j\in\mathcal{J}_\emptyset} \left[v(\theta_t, j) - p_{j,t}\right]\xi_t + \sum_{j\in\mathcal{J}} p_{j,t} d_{j,t} \mid \mathbf{q}_t\Big] \\
&= \mathbb{E}[\xi_t]\mathbb{E}\Big[\max_{j\in\mathcal{J}_\emptyset} \left[v(\theta_t, j) - p_{j,t}\right] \mid \mathbf{q}_t\Big] + \mathbb{E}\Big[\sum_{j\in\mathcal{J}} p_{j,t} d_{j,t} \mid \mathbf{q}_t\Big] \qquad (11) \\
&= \frac{\lambda}{1+\lambda}\int_\Theta \max_{j\in\mathcal{J}_\emptyset} \left[v(\theta_t, j) - p_{j,t}\right] dF(\theta) + \frac{1}{1+\lambda}\sum_{j\in\mathcal{J}} \mu_j p_{j,t} \\
&= \frac{\lambda}{1+\lambda} h(\mathbf{p}_t)\,.
\end{aligned}
$$

By adding $c \cdot v(\theta_t, a(\theta_t, \mathbf{q}_t))\xi_t$ to both sides of equation (10), taking expectation conditional on $\mathbf{q}_t$ and applying equation (11) we obtain the required identity. □

Lemma 4. *If $c(w) = c \cdot w$, we have that for any $\mathbf{q}_t$,*

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

Proof. We have

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

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

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

□

Lemma 5. *For the model defined in Section 3, we have that*

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

Proof. We need to argue that the limiting operator and the expectation can be interchanged. We proceed in 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,[12] 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 limit and expectation. It follows from the 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 non-negative 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 implies 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

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

we have

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

□

[12] 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.

4.1 Asymptotic Optimality

Theorem 1 bounds the allocate efficiency loss in the waiting list. The following corollaries imply that the loss tends to zero when agents' patience increases or when the market size grows large.

As agents are increasingly patient, the marginal expected waiting cost becomes smaller. For example if $c(w) = c \cdot w$, then $\Delta = c/\mu_{\min} \to 0$ as $c \to 0$. This implies:

Corollary 1. *Fix $F, \{\mu\}_{j\in\mathcal{J}}, \lambda$ and consider a sequence of markets indexed by ℓ in which waiting cost is $c(w) = c_\ell \cdot w$. Let W_ℓ^{WL} denote the allocative efficiency of the waiting list for market ℓ. If $c_\ell \to 0$ as $\ell \to \infty$, then*

$$W_\ell^{\mathrm{WL}} \xrightarrow[\ell\to\infty]{} W^{\mathrm{OPT}}\,.$$

Note that W^{OPT} and the optimal prices $\mathbf{p}^*$ are independent of c_ℓ. As c_ℓ tends to zero, the length of queue j that generates the optimal price, $q_j^* = p_j^*/c_\ell$, is increasing. Moreover, there are smaller fluctuations in queue lengths and there fore in prices.This implies lower variation in prices and higher allocative efficiency.

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 lower. For example, if $c(w) = c \cdot w$ and $\mu_{\min} \to \infty$, then $\Delta = c/\mu_{\min} \to 0$. Therefore:

Corollary 2. *Fix $F, \{\mu\}_{j\in\mathcal{J}}, \lambda$ and $c(\cdot)$. Consider a sequence of markets indexed by ℓ in the agent arrival rate is $\ell \cdot \lambda$ and the arrival rate of item j is $\ell \cdot \mu_j$. Let W_ℓ^{WL} denote the allocative efficiency of the waiting list for market ℓ. If $\ell \to \infty$ then*

$$W_\ell^{\mathrm{WL}} \xrightarrow[\ell\to\infty]{} W^{\mathrm{OPT}}\,.$$

4.2 Lower Bound for the Allocative Efficiency Loss

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

Example 1. *Consider a dynamic 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 agent type. The distribution of agent types and 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*

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

Proposition 2. *For any J and any $\Delta > 0$ and $\varepsilon > 0$, there exists a market with J items as in Example 1 in which the allocative efficiency loss under the waiting list is*

$$W^{\mathrm{OPT}} - W^{\mathrm{WL}} \geq \Delta - \varepsilon$$

Proof of Proposition 2. We prove the result by calculating the allocative efficiency loss in the market of Example 1. 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 $j = \theta$. An agent arriving in epoch t of type $\theta_t = j$ will choose to join queue j to receive a value of 1 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 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 $\pi_j(k) = \frac{1}{K+1}$.[13]

The allocative efficiency under the waiting list is given by

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

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

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

□

[13] 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 $\pi_j(0) = \pi_j(1) = \cdots = \pi_j(K)$.

4.3 Welfare of the Queueing Mechanism

As the market grows large, the welfare converges to the optimal allocation minus the predicted waiting times. This follows from the following result that restricts attention to the convergence of prices.

Proposition 3. *Fix $F, \{\mu\}_{j\in\mathcal{J}}, \lambda$ and consider a sequence of markets indexed by ℓ in which waiting cost is $c(w) = c_\ell \cdot w$. Let $\mathcal{P}^*$ be the set of minimizers of the dual function $h(\mathbf{p})$, let $\mathbf{p}_{\ell,\infty}$ be the random variable of steady-state price in the market indexed by ℓ. If $c_\ell \to 0$ as $\ell \to \infty$, then*

$$\limsup_{\ell\to\infty} \ \mathbb{P}\left(\mathbf{p}_{\ell,\infty} \notin \mathcal{P}^*\right) = \ 0\,.$$

Proof sketch. Let $\mathbf{p}^*$ be any minimizer of the dual function $h(\mathbf{p})$. Consider the following centered Lyapunov function:

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

We first fix ℓ and omit the subscript ℓ. Recall that $p_{j,t+1} = p_{j,t} + \frac{c}{\mu_j}(a_{j,t} - d_{j,t} + u_{j,t})$. Similar to the proof of Lemma 3, we have

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

Here inequality (a) is because $p_j^* \geq 0$ and $u_{j,t} \geq 0$. Note that $(p_{j,t} + \frac{c}{\mu_j}(a_{j,t} - d_{j,t}))u_{j,t} \equiv 0$. Using arguments similar to those in the proof of Lemma 4, we have

$$\bar{L}(\mathbf{p}_{t+1}) - \bar{L}(\mathbf{p}_t) \leq \ c\langle \mathbf{p}_t - \mathbf{p}^*, \mathbf{a}_t - \mathbf{d}_t\rangle + c\cdot\Delta\,.$$

Take expectations on both sides conditioned on $\mathbf{q}_t$, and use the fact that $\mathbb{E}[\mathbf{a}_t - \mathbf{d}_t|\mathbf{q}_t] \in$

$-\partial h(\mathbf{p}_t)$, we have

$$\mathbb{E}[\bar{L}(\mathbf{p}_{t+1})|\mathbf{q}_t] - \bar{L}(\mathbf{p}_t) \leq \ -c\,(h(\mathbf{p}_t) - h(\mathbf{p}^*)) + c \cdot \Delta\,. \tag{12}$$

For $\epsilon > 0$, let $B(\epsilon) \triangleq \max\{\bar{L}(\mathbf{p}) : h(\mathbf{p}) \leq h(\mathbf{p}^*) + \epsilon\}$. It is straightforward to verify that for the Markov chain $\{\mathbf{p}_t\}_{t=1}^{\infty}$ and $\Delta \leq \frac{\epsilon}{2}$, the Lyapunov function $\bar{L}(\cdot)$ has negative drift $-\frac{c\epsilon}{2}$ when $h(\mathbf{p}) \geq h(\mathbf{p}^*) + \epsilon$. Notice that the maximum increase of $\bar{L}(\mathbf{p})$ in one period is $cv_{\max}$. Applying Theorem 1 in Bertsimas et al. (2001), we have

$$\mathbb{P}(\bar{L}(\mathbf{p}_\infty) > B(\epsilon) + 2cv_{\max}m) \leq \left(\frac{2v_{\max}}{2v_{\max} + \epsilon}\right)^{m+1}$$

Let $m = \frac{B(\epsilon)}{cv_{\max}}$. Let $\ell \to \infty$ therefore $c_\ell \to 0$, and this shows that $\mathbb{P}(\bar{L}(\mathbf{p}_{\ell,\infty}) > 3B(\epsilon)) \to 0$. Since ϵ can be chosen arbitrarily, this concludes the proof. □

5 Exponentially Small Loss for Generic Problems

The magnitude of the loss generated by miss-allocations depends naturally on the parameters of the economy and this section further isolates sources of the inefficiency. Indeed it is shown here that in economies with finitely many agents and item types, the loss is generically much smaller than the stated in Theorem 1. The following natural assumption will drive the result.

Assumption 1 (Unique shadow price)**.** *The dual problem* (5) *with finite types of items and agents has a unique minimizer* $\mathbf{p}^*$*.*

Assumption 1 holds generically when there are finitely many agent types.

Proposition 4. *Let* $\boldsymbol{\lambda} \in \mathbb{R}_{++}^{|\Theta|}$ *be the vector of arrival rates of each type of agent. The vectors* $(\mathbf{v}, \boldsymbol{\lambda}, \boldsymbol{\mu})$ *satisfying Assumption 1 are open and dense in* $\mathbb{R}_{++}^{|\Theta|\times|\mathcal{J}|+|\Theta|+|\mathcal{J}|}$*.*

Proof of Proposition 4. We say that the problem instance $(\boldsymbol{\lambda}, \boldsymbol{\mu})$ satisfies *generalized imbalance (GI)* if there are 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 $\mathcal{J}'$. The proposition holds because the problem instances satisfying GI are open and dense in $\mathbb{R}_{++}^{|\mathcal{I}|+|\mathcal{J}|}$, and that GI implies Assumption 1 (see Proposition C.2 of Johari et al. (2016)). □

Theorem 2. *Suppose there are a finite number of agent types, the primitives* $\mathbf{v} = (v(\theta, j))_{\theta\in\Theta, j\in\mathcal{J}}$*,* $\boldsymbol{\lambda} = (\lambda_\theta)_{\theta\in\Theta}$*,* $\boldsymbol{\mu} = (\mu_j)_{j\in\mathcal{J}}$ *satisfy Assumption 1, and consider a sequence of markets indexed by* ℓ *in which waiting cost is* $c(w) = c_\ell \cdot w$*. If* $c_\ell \to 0$ *as* $\ell \to \infty$*, then there exists*

$\alpha = \alpha(\mathbf{v}, \boldsymbol{\lambda}, \boldsymbol{\mu}) > 0$, *such that*

$$W^{\mathrm{OPT}} - W_\ell^{\mathrm{WL}} = O\left(\exp(-\alpha/\Delta_\ell)\right) \quad \ell \to \infty\,.$$

The intuition behind Theorem 2 is the following. Assumption 1 led to the robustness of the optimal dual prices; indeed as long as prices are within some $\delta > 0$ distance from the unique dual prices $\mathbf{p}^*$, items are allocated efficiently due to complementary slackness. And when prices are further away from $\mathbf{p}^*$, they adjust towards the optimal dual prices at some minimal rate $\gamma > 0$. So prices follow a biased random walk towards $\mathbf{p}^*$, implying that that they are far from the optimal dual prices with probability that is exponentially small in step size Δ_ℓ's inverse. The exponent α in the theorem is roughly proportional to the product of δ and γ.

The remainder of this section formalizes this intuition. For this purpose consider an economy that satisfies Assumption 1 with the unique dual price $\mathbf{p}^*$ and assume that agents have linear waiting costs as stated in the theorem.

Robustness of dual prices. The following sets will serve to define the robustness of dual prices. Define the set of *active agent types* as:

$$\Theta^* \triangleq \left\{\theta \in \Theta : \max_{j \in \mathcal{J}} \left(v(\theta, j) - p_j^*\right) > 0\right\},$$

the set of *active item types* as:

$$\mathcal{J}^* \triangleq \left\{j \in \mathcal{J} : p_j^* > 0\right\},$$

and for each type of agent $\theta \in \Theta$, its set of *active matches* as:

$$\mathcal{J}_\theta^* \triangleq \operatorname{argmax}_{j \in \mathcal{J}} [v(\theta, j) - p_j^*]^+\,.$$

From complementary slackness conditions, an allocation $\mathbf{x}$ is an optimal solution of the static allocation problem (4) *if and only if* (i) all active agent types are assigned items, (ii) all active item types are assigned to agents, (iii) for each (θ, j) such that $x_{\theta,j} > 0$, (θ, j) is an active match, and, (iv) $\mathbf{x}$ satisfies the primal constraints.

For uniquely defined $\mathbf{p}^*$, observe that there exists a constant $\delta > 0$ such that when the price is within δ of $\mathbf{p}^*$, conditions (i)-(iii) are satisfied.

Definition 2. *Consider a problem instance* $(\mathbf{v}, \boldsymbol{\lambda}, \boldsymbol{\mu})$ *that satisfies Assumption 1. Let*

$$\delta_{\boldsymbol{\lambda}} \triangleq \min_{\theta \in \Theta^*} \max_{j \in \mathcal{J}} \left(v(\theta, j) - p_j^*\right),$$

$$\delta_{\boldsymbol{\mu}} \triangleq \min_{j \in \mathcal{J}^*} \; p^*_j \,,$$
$$\delta_{\mathbf{v}} \triangleq \min_{\{\theta \in \Theta : \mathcal{J}^*_\theta \neq \emptyset\}} \; \min_{j \in \mathcal{J}^*_\theta, j' \notin \mathcal{J}^*_\theta} \left([v(\theta, j) - p^*_j]^+ - v(\theta, j') + p^*_{j'} \right) .$$

We refer to $\delta \triangleq \min\{\delta_{\boldsymbol{\lambda}}, \delta_{\boldsymbol{\mu}}, \delta_{\mathbf{v}}\} > 0$ *as the* robustness of dual price *of the problem instance.*

Define the following set of near-optimal dual prices.

$$\mathcal{P} \triangleq \left\{ \mathbf{p} : ||\mathbf{p} - \mathbf{p}^*||_\infty < \frac{\delta}{2} \right\} .$$

The following lemma and its proof states that the loss arises only when $\mathbf{p} \notin \mathcal{P}$.

Lemma 6. *If prices in steady state satisfy that* $\mathbb{P}(\mathbf{p}_\infty \notin \mathcal{P}) \leq \kappa$*, then*

$$W^{\mathrm{OPT}} - W^{\mathrm{WL}} \leq \; \frac{1+\lambda}{\lambda} v_{\max} \kappa \,.$$

Proof. Let $x_{\theta,j,t}$ be the match in the t-th period, which equals to 1 if a type θ agent is assigned a type j item in the t-th period, and equals to zero otherwise. For prices that satisfy $\mathbf{p}_t \in \mathcal{P}$, it holds by definition (robustness of the dual price) that

$$\sum_{j \in \mathcal{J}} x_{\theta,j,t} = \; \mathbb{1}\{\text{a type } \theta \text{ agent arrives at } t\} \,, \qquad \forall \theta \in \Theta^*$$
$$\sum_{\theta \in \Theta} x_{\theta,j,t} = \; \mathbb{1}\{\text{a type } j \text{ item arrives at } t\} \,, \qquad \forall j \in \mathcal{J}^*$$
$$x_{\theta,j,t} = 0 \,, \qquad \forall \theta \in \Theta \,, \theta \notin \mathcal{J}^*_\theta \,.$$

Denote by w_∞ the match value obtained in a period in steady state (by the Arrival Theorem, the steady state of the discrete time process we are considering is the same as the steady state distribution of the original continuous time process). Note that

$$W^{\mathrm{WL}} = \; \frac{1+\lambda}{\lambda} \mathbb{E}[w_\infty] \geq \; \frac{1+\lambda}{\lambda} \mathbb{E}[w_\infty | \mathbf{p}_\infty \in \mathcal{P}] \cdot \mathbb{P}(\mathbf{p}_\infty \in \mathcal{P}) \,. \tag{13}$$

The last term on the RHS is bounded from below $1 - \kappa$ by assumption. It remains to show $\frac{1+\lambda}{\lambda} \mathbb{E}[w_\infty | \mathbf{p}_\infty \in \mathcal{P}] = W^{\mathrm{OPT}}$ because

$$W^{\mathrm{WL}} \geq \; (1-\kappa) W^{\mathrm{OPT}} \geq W^{\mathrm{OPT}} - \frac{1+\lambda}{\lambda} v_{\max} \kappa \,.$$

By linearity of expectation, we have

$$\mathbb{E}[w_\infty | \mathbf{p}_\infty \in \mathcal{P}] = \sum_{\theta \in \Theta, j \in \mathcal{J}} v(\theta, j) \mathbb{E}[x_{\theta,j,\infty} | \mathbf{p}_\infty \in \mathcal{P}] \,.$$

Note that

$$\sum_{j\in\mathcal{J}} \mathbb{E}[x_{\theta,j,\infty}|\mathbf{p}_\infty \in \mathcal{P}] = \lambda_\theta\,, \qquad \forall \theta \in \Theta^* \tag{14}$$

$$\sum_{\theta\in\Theta} \mathbb{E}[x_{\theta,j,\infty}|\mathbf{p}_\infty \in \mathcal{P}] = \mu_j\,, \qquad \forall j \in \mathcal{J}^* \tag{15}$$

$$\mathbb{E}[x_{\theta,j,\infty}|\mathbf{p}_\infty \in \mathcal{P}] = 0\,, \qquad \forall \theta \in \Theta\,, j \notin \mathcal{J}_\theta^*\,. \tag{16}$$

This completes the proof since equalities (14)-(16) correspond to the complementary slackness condition of the static allocation problem. □

Rate of adjustment. Next we show that prices are very likely to be in $\mathcal{P}$, by showing that prices that deviates from $\mathbf{p}^*$ quickly adjusts back. The rate of price adjustment is related to the "sharpness" of the dual objective (5).

Lemma 7 (Geometry of dual function)**.** *Suppose Assumption 1 holds. Then there exists* $\gamma(\mathbf{v},\boldsymbol{\lambda},\boldsymbol{\mu}) > 0$ *such that for any* $\mathbf{p} \in \mathbb{R}^{|\mathcal{J}|}$*, we have*

$$h(\mathbf{p}) - h(\mathbf{p}^*) \geq \gamma(\mathbf{v},\boldsymbol{\lambda},\boldsymbol{\mu})||\mathbf{p}-\mathbf{p}^*||_2\,. \tag{17}$$

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 of this technical lemma appear in Appendix D.1. The proof shows that the rate of adjustment is positive for $\mathbf{p} \in \mathcal{P}$ and by convexity of $h(\mathbf{p}) - h(\mathbf{p}^*)$ this holds also also for prices not in $\mathcal{P}$. The proof Theorem 2 can now be completed.

[PQ: current progress]

Proof of Theorem 2. Using Lemma 6, we have that

$$W^{\mathrm{OPT}} - W^{\mathrm{WL}} \leq \frac{1+\lambda}{\lambda} v_{\max}\mathbb{P}(\mathbf{p}_\infty \notin \mathcal{P})\,.$$

Lemma 12 in the appendix provides the following concentration bound (using further Lyapunov analysis and using the established bound on the rate of adjustment):

$$\mathbb{P}(\mathbf{p}_\infty \notin \mathcal{P}) \leq \exp\left(-\log\left(1+\frac{\gamma\mu_{\min}}{4}\right)\left(\frac{\delta}{12\Delta}\right)\right).$$

Therefore the allocative efficiency loss is bounded by:

$$W^{\mathrm{OPT}} - W^{\mathrm{WL}} \leq \frac{1+\lambda}{\lambda} v_{\max} \exp\left(-\log\left(1+\frac{\gamma\mu_{\min}}{4}\right)\left(\frac{\delta}{12\Delta}\right)\right).$$

where δ is the robustness of price, and γ is the rate of price adjustment. □

6 On the Optimal Step Size

In our main results we considered a waiting list, for which the step-size Δ is exogenously given by the item arrival rates and agents' patience. In this section we enrich the model to consider the optimal step-size for a planner who implements a price adaptation to maximize allocative efficiency. We assume that the planner can fix a single SGD step size in order to minimize the loss due price fluctuations. The model assumes linear waiting costs and finitely many periods T. This section illustrates that for a carefully chosen step size, the loss can grow at a rate of $\sqrt{T}$.

Assignments and Allocative Efficiency in Finite Horizon T. Similar to (1), we define the allocative efficiency of allocation η in the first T epochs as follows:

$$W_T(\eta) = \sum_{t=1}^{T} \xi_t v(\theta_t, \eta_t). \tag{18}$$

Note that for the finite horizon setting, it is notationally easier to work with the sum of match values rather than the average match values, therefore we use the sum as objective. A caveat is that since we are only interested in the first T epochs, we can only assign items that arrive before T to the agents.

Define optimal allocative efficiency to be

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

where H is the set of all allocations of interest and the expectation is taken over all possible realizations.

Proposition 5. *The optimal finite horizon allocative efficiency W_T^{OPT} is upper bounded by $W_T^* \triangleq \frac{\lambda}{\lambda+1} \cdot T \cdot W^*$, where W^* is the optimal static allocative effciency (i.e., the optimal value of (4)).*

The proof of Proposition (5) is almost identical to the proof of Proposition (1) and is therefore omitted.

We first show that by setting the step size to $1/\sqrt{T}$, we can obtain allocative efficiency that is within $O(\sqrt{T})$ of W_T^*, hence also within $O(\sqrt{T})$ of W_T^{OPT}.

Theorem 3. *The T-horizon allocative efficiency under the SGD algorithm with step size $\Delta = 1/\sqrt{T}$ satisfies:*

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

Proof. From the analysis of Theorem 1, equation (9), we have:

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

Therefore the allocative efficiency loss satisfies:

$$W_T^* - W_T^{\mathrm{WL}} = \frac{1}{c}\left(L(\mathbf{p}_1) - \mathbb{E}[L(\mathbf{p}_{T+1})]\right) + T\frac{2+\lambda}{2(1+\lambda)}\Delta\,.$$

Note that $\Delta = c/\mu_{\min}$, and that

$$\frac{1}{\mu_{\min}\cdot\Delta}\left(L(\mathbf{p}_1) - \mathbb{E}[L(\mathbf{p}_{T+1})]\right) + T\frac{2+\lambda}{2(1+\lambda)}\Delta \leq \ T\frac{2+\lambda}{2(1+\lambda)}\Delta + \frac{1}{2\mu_{\min}\cdot\Delta}\sum_{j\in\mathcal{J}}\mu_j p_{1,j}^2\,.$$

By setting $\Delta = 1/\sqrt{T}$, the loss is upper bounded by

$$\sqrt{T}\left(\frac{2+\lambda}{2(1+\lambda)} + \frac{1}{2\mu_{\min}} v_{\max}^2\right).$$

This concludes the proof. □

We show that the bound given in Theorem 3 is tight in the sense that there exists an economy (which we explicitly construct below) for which $W_T^* - W_T^{\mathrm{OPT}} = \Omega(\sqrt{T})$.

Example 2. *Consider a dynamic market with a single type of item which arrives according to a Poisson process with rate* 1*, and a single type of agent which also arrives according to a Poisson process with rate* 1*. The agent's valuation of item is* 1*.*

Proposition 6. *For the market in Example 2, we have*

$$W_T^* - W_T^{\mathrm{OPT}} \geq 0.2\sqrt{T} - 5\,.$$

Proof. It is straightforward to derive that $W_T^* = T$.

On the other hand, note that the optimal assignment simply waits until the last period to match all the agents and items that have arrived. Let $\Lambda(T)$, $\Lambda'(T)$ be two Poisson random

variables with mean T. We have

$$W_T^{\mathrm{OPT}} = \mathbb{E}[\min\{\Lambda(T), \Lambda'(T)\}]\,.$$

Therefore

$$\begin{aligned} W_T^* - W_T^{\mathrm{OPT}} &= T - \mathbb{E}[\Lambda'(T) + \min\{\Lambda(T) - \Lambda'(T), 0\}] \\ &= \mathbb{E}[\max\{\Lambda'(T) - \Lambda(T), 0\}] \\ &= \sqrt{2T}\ \mathbb{E}\left[\max\left\{\frac{\Lambda'(T) - \Lambda(T)}{\sqrt{2T}}, 0\right\}\right]\,. \end{aligned}$$

By Markov's inequality, if follows that

$$\mathbb{E}\left[\max\left\{\frac{\Lambda'(T) - \Lambda(T)}{\sqrt{2T}}, 0\right\}\right] \geq \mathbb{P}\left(\max\left\{\frac{\Lambda'(T) - \Lambda(T)}{\sqrt{2T}}, 0\right\} \geq 1\right)\,.$$

Therefore we can write that

$$W_T^* - W_T^{\mathrm{OPT}} \geq \sqrt{2T}(1 - F_T(1)) \tag{19}$$

where $F_T(x)$ is the cumulative distribution function of $\max\left\{\frac{\Lambda'(T)-\Lambda(T)}{\sqrt{2T}}, 0\right\}$. Let $\Phi(x)$ be the cumulative distribution function of standard Gaussion distribution, we have

$$1 - F_T(1) = 1 - \Phi(1) + \Phi(1) - F_T(1) \geq 1 - \Phi(1) - |\Phi(1) - F_T(1)|\,.$$

We can view $\Lambda'(T) - \Lambda(T)$ as the sum of T i.i.d. random variables $\Lambda'(1) - \Lambda(1)$, where $\Lambda(1), \Lambda'(1)$ are independent Poisson distribution with mean 1. Using Berry-Esseen bound we have

$$|\Phi(1) - F_T(1)| \leq \frac{0.4748\ \mathbb{E}|\Lambda(1) - \Lambda'(1)|^3}{[\mathrm{Var}(\Lambda'(1) - \Lambda(1))]^{3/2}\sqrt{T}}\,. \tag{20}$$

Note that $\mathrm{Var}(\Lambda'(1) - \Lambda(1)) = 2$, and by Esseen (1975) we have

$$\mathbb{E}|\Lambda(1) - \Lambda'(1)|^3 \leq 2\mathbb{E}[|\Lambda(1)|^3] + 3\mathbb{E}[\Lambda(1)]\mathbb{E}[(\Lambda(1))^2] - \mathbb{E}^3[\Lambda(1)] = 15\,.$$

Plugging into (20), we have

$$|\Phi(1) - F_T(1)| \leq \frac{3}{\sqrt{T}}\,.$$

Plugging back into (19) and note that $1 - \Phi(1) \geq 0.1587$, we have

$$W_T^* - W_T^{\mathrm{OPT}} \geq \sqrt{2T}(1 - \Phi(1) - 3/\sqrt{T}) \geq 0.2\sqrt{T} - 5\,.$$

This concludes the proof. □

7 Summary

This paper considers a dynamic economy, in which the waiting times play the role of prices in guiding the allocation and rationing items. It studies the impact of the fluctuations in waiting times resulting from the stochasticity in the arrival of demand and supply, and quantifies the allocative efficiency loss resulting from this fluctuation.

We observe that the efficiency loss compared to an optimum offline assignment is fairly small and is captured by the marginal increase in expected waiting cost from having one more agent in the queue. Furthermore, when equilibrium waiting times are essentially unique, there is almost no efficiency loss.

Our results show that despite the decentralized nature of these markets and the underlying stochasticity, simple waiting mechanisms obtain near optimum allocative efficiency. They also justify deterministic or "fluid" approximations for modeling the behavior of such queueing systems.

References

Afeche, P., Caldentey, R., and Gupta, V. (2019). On the optimal design of a bipartite matching queueing system. *Available at SSRN 3345302*.

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

Ata, B., Ding, Y., and Zenios, S. (2021). An achievable-region-based approach for kidney allocation policy design with endogenous patient choice. *Manufacturing & Service Operations Management*, 23(1):36–54.

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

BackProp, E., LeCun, Y., Bottou, L., Orr, G. B., and Muller, K. (1998). Neural networks: Tricks of the trade.

Bertsekas, D. P. (1979). A distributed algorithm for the assignment problem. *Lab. for Information and Decision Systems Working Paper, MIT*.

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.

Braverman, M., Chen, J., and Kannan, S. (2016). Optimal provision-after-wait in healthcare. *Mathematics of Operations Research*, 41(1):352–376.

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.

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.

Esseen, C.-G. (1975). Bounds for the absolute third moment. *Scandinavian Journal of Statistics*, pages 149–152.

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*.

Georgiadis, L., Neely, M. J., Tassiulas, L., et al. (2006). Resource allocation and cross-layer control in wireless networks. *Foundations and Trends® in Networking*, 1(1):1–144.

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.

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.

Polterovich, V. (1993). Rationing, queues, and black markets. *Econometrica: Journal of the Econometric Society*, pages 1–28.

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.

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.

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.

Organization of the Mathematical Appendices. The appendix is organized as follows.

1. In Appendix A, we prove Proposition 1, which shows that the optimal allocative efficiency equals the value of the static allocation problem.
2. In Appendix B, we prove Theorem 1.
3. In Appendix C, we bound the price change granularity Δ using the property of the waiting cost functions.
4. In Appendix D, we establish the genericity of Assumption 1, and prove Theorem 2, which shows that for generic instances with finite agent types, the efficiency loss of the queueing mechanism is exponentially small in N as market size N increases.

A Optimal Allocative Efficiency: Proof of Proposition 1

In this section we prove Proposition 1.

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

Let $\eta \in H$ be any no-Ponzi allocation. Recall that $\eta_t \in \mathcal{J}_\emptyset$ is the kind of item assigned under η to the agent that arrived at epoch t. For each $j \in \mathcal{J}$ and $\theta \in \Theta$, define $\hat{G}_j^T(\theta)$ as

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

Recall that A_T is the number of agents that arrived in the first T epochs. Therefore, $\hat{G}_j^T(\theta)$ is proportional to the empirical cumulative distribution function of the types of the agents in A_T who 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

$$W(\eta) = \liminf_{T\to\infty} \frac{1}{A_T} \sum_{j\in\mathcal{J}} \sum_{t\leq T} \xi_t v(\theta_t, j) = \liminf_{T\to\infty} \sum_{j\in\mathcal{J}} \int_{\theta\in\Theta} v(\theta, j) d\hat{G}_j^T(\theta) . \tag{21}$$

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

$$\sum_{j\in\mathcal{J}} \hat{G}_j^T(\theta) \leq \frac{1}{A_T} \sum_{t\leq T} \xi_t \mathbb{1}_{\{\theta_t \leq \theta\}} , \qquad \forall \theta \in \Theta \tag{22}$$

$$\hat{G}_j^T(0) = 0 , \hat{G}_j^T(1) \leq \frac{1}{A_T} \left(\sum_{t\leq T} \mathbb{1}_{\{\eta_t = j\}} + M \right) \qquad \forall j \in \mathcal{J} \tag{23}$$

$$\hat{G}_j^T(\theta) \text{ is non-decreasing and right-continuous.} \qquad \forall j \in \mathcal{J} \qquad (24)$$

for some $M \in \mathbb{R}$. Here (22) and (24) are trivial. (23) is satisfied by any no-Ponzi assignment for the following reason: The agents in 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 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

$$\mathbb{E}\left[\sum_{j\in\mathcal{J}} \int_{\theta\in\Theta} v(\theta, j) d\hat{G}_j^T(\theta)\right] \le \mathbb{E}\left[\max_{\hat{G}_j(\theta) \text{ satisfying (22)(23)(24)}} \sum_{j\in\mathcal{J}} \int_{\theta\in\Theta} v(\theta, j) d\hat{G}_j(\theta)\right].$$

It is easy to check that the optimal value of the inner maximization problem above is concave and non-decreasing in the RHS of (22) and (23). Note that

$$\begin{aligned} &\text{Expectation of RHS of (22)} = F(\theta), && \forall \theta \in \Theta, \\ &\text{Expectation of RHS of (23)} = \frac{\mu_j}{\lambda} + \frac{(1+\lambda)M}{\lambda T}, && \forall j \in \mathcal{J}. \end{aligned}$$

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 (4). Therefore $\mathbb{E}[W(\eta)] \le W^*$ for any $\eta \in H$. Since $W(\eta)$ is uniformly bounded above by $v_{\max}$, by 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 concludes the proof.

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

Denote the optimal solution of the optimization problem (4) 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 there is any; and the item is discarded if the item's queue is empty. The coin-toss is defined as follows: Fix $M \in \mathbb{Z}_+$. If the arriving agent is of type θ, it is assigned to queue j with probability $x^*_{\theta j}$, or rejected 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 as:

$$\begin{aligned} W^* - W^{\mathrm{RD}} &\leq W^* - W^* \cdot \mathbb{P}(q_{j,\infty} < M, \forall j \in \mathcal{J}) \\ &\leq v_{\max} \cdot \mathbb{P}(q_{j,\infty} = M, \exists j \in \mathcal{J}) \\ &\leq v_{\max} \cdot \sum_{j\in\mathcal{J}} \mathbb{P}(q_{j,\infty} = M) \,, \end{aligned} \tag{25}$$

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 (4) 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 (25), we have

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

Notice that by definition, $W^{\mathrm{RD}} \leq W^{\mathrm{OPT}}$, 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 concludes the proof. □

B Upper Bound on the Allocative Efficiency Loss of the Queueing Mechanism

In this section we prove Theorem 1.

Proof of Theorem 1. The following proof generalizes the one described in the main paper which focuses on the special case of linear waiting cost.

Recall that $W_T(\eta^{\mathrm{WL}})$ is the total value of items assigned to agents that arrive before epoch T, that is

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

and that

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

Similar to the proof in the main paper, we use the Lyapunov analysis to bound the allocative efficiency. It turns out for general waiting costs, using the Lyapunov function of *queue lengths* is notationally simpler than using the Lyapunov function of *waiting costs*. As a result, we use the following Lyapunov function: let the Lyapunov function $L(\mathbf{q})$ be 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 follows

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

Let $\mathbf{a}_t$ and $\mathbf{d}_t$ be the vectors representing the arriving agent and item at time t, respectively:

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

and let $u_{j,t} \triangleq \max\{0, d_{j,t} - q_{j,t} - a_{j,t}\}$ denote the number of discarded items of type j at time t. The evolution of the length of queue j is governed by

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

By Lemma 9 we have that

$$\begin{aligned}\mathbb{E}[v(\theta_t, a(\theta_t, \mathbf{q}_t))\xi_t|\mathbf{q}_t] \geq \frac{\lambda}{1+\lambda}h(\mathbf{p}_t) - (L(\mathbf{q}_t) - \mathbb{E}[L(\mathbf{q}_{t+1}) \mid \mathbf{q}_t]) - \frac{\Delta}{2(1+\lambda)} \\ - \mathbb{E}[D_L(\mathbf{q}_t + \mathbf{a}_t - \mathbf{d}_t, \mathbf{q}_t) \mid \mathbf{q}_t]\,.\end{aligned} \tag{26}$$

By Lemma 10 we have that $D_L(\mathbf{q}_t + \mathbf{a}_t - \mathbf{d}_t, \mathbf{q}_t) \leq \Delta/2$. By Lemma 1 we have that $h(\mathbf{p}_t) \geq W^*$. Together, we have that

$$\mathbb{E}[v(\theta_t, a(\theta_t, \mathbf{q}_t))\xi_t|\mathbf{q}_t] \geq \frac{\lambda}{1+\lambda}W^* - (L(\mathbf{q}_t) - \mathbb{E}[L(\mathbf{q}_{t+1}) \mid \mathbf{q}_t]) - \left(\frac{\Delta}{2(1+\lambda)} + \frac{\Delta}{2}\right).$$

Therefore, we have that

$$\begin{aligned}\mathbb{E}\left[W_T(\eta^{\mathrm{WL}})\right] &= \mathbb{E}\left[\sum_{t=1}^{T} \xi_t \cdot v(\theta_t, a(\theta_t, \mathbf{q}_t))\right] \\ &= \mathbb{E}\left[\mathbb{E}\left[\sum_{t=1}^{T} \xi_t \cdot v(\theta_t, a(\theta_t, \mathbf{q}_t)) \mid \mathbf{q}_t\right]\right] \\ &= \mathbb{E}\left[\sum_{t=1}^{T} \mathbb{E}\left[\xi_t \cdot v(\theta_t, a(\theta_t, \mathbf{q}_t)) \mid \mathbf{q}_t\right]\right] \\ &\geq \mathbb{E}\left[\sum_{t=1}^{T} \frac{\lambda}{1+\lambda}W^* - (L(\mathbf{p}_t) - \mathbb{E}[L(\mathbf{q}_{t+1}) \mid \mathbf{q}_t]) - \frac{2+\lambda}{2(1+\lambda)}\Delta\right] \\ &= T\frac{\lambda}{1+\lambda}W^* - (L(\mathbf{p}_1) - \mathbb{E}[L(\mathbf{q}_{T+1})]) - T\frac{2+\lambda}{2(1+\lambda)}\Delta\,.\end{aligned} \tag{27}$$

By Lemma 5, we have

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

Plugging in (27) to the above equality, we have

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

This concludes the proof. □

Lemma 8. *We have that*

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

Proof. By definition of Bregman divergence, we have

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

Therefore

$$L(\mathbf{q}_{t+1}) = \; L(\mathbf{q}_t + \mathbf{a}_t - \mathbf{d}_t) + \langle \mathbf{p}(\mathbf{q}_{t+1}), \mathbf{u}_t \rangle - D_L(\mathbf{q}_t + \mathbf{a}_t - \mathbf{d}_t, \mathbf{q}_{t+1})\,. \tag{28}$$

To bound the RHS of (28), 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 $u_{j,t} = 1$. 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

$$\begin{aligned} &\langle \mathbf{p}(\mathbf{q}_{t+1}), \mathbf{u}_t \rangle - D_L(\mathbf{q}_t + \mathbf{a}_t - \mathbf{d}_t, \mathbf{q}_{t+1}) \\ &= p_j(0) - (P_j(-1) - P_j(0) - p_j(0) \cdot (-1)) \\ &= P_j(0) - P_j(-1)\,. \end{aligned} \tag{29}$$

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

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

Plugging in the above equality to (29), we have

$$\langle \mathbf{p}(\mathbf{q}_{t+1}), \mathbf{u}_t \rangle - D_L(\mathbf{q}_t + \mathbf{a}_t - \mathbf{d}_t, \mathbf{q}_{t+1}) \le \frac{\Delta}{2} \cdot \mathbb{1}\{\xi_t = 0\}\,. \tag{30}$$

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

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

Therefore, plugging in the above two cases to (28), we have

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

□

Lemma 9. *For the model defined in Section 3, we have that*

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

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

$$\begin{aligned} & L(\mathbf{q}_t) - L(\mathbf{q}_{t+1}) \\ \geq\ & L(\mathbf{q}_t) - L(\mathbf{q}_t + \mathbf{a}_t - \mathbf{d}_t) - \frac{\Delta}{2} \cdot \mathbb{1}\{\xi_t = 0\} \\ =\ & -\langle \mathbf{p}(\mathbf{q}_t), \mathbf{a}_t - \mathbf{d}_t \rangle - D_L(\mathbf{q}_t + \mathbf{a}_t - \mathbf{d}_t, \mathbf{q}_t) - \frac{\Delta}{2} \cdot \mathbb{1}\{\xi_t = 0\} , \end{aligned} \tag{31}$$

where the inequality follows from Lemma 8, and the equality comes from the definition of Bregman divergence.

We expanded the expected value of the next arrival $\mathbb{E}\Big[v(\theta_t, a(\theta_t, \mathbf{q}_t))\xi_t \mid \mathbf{q}_t\Big]$ plus the expected value of the term $\langle \mathbf{p}_t, \mathbf{a}_t - \mathbf{d}_t \rangle$ to show these are related to the dual objective given in (5).

$$\begin{aligned} & \mathbb{E}\Big[v(\theta_t, a(\theta_t, \mathbf{q}_t))\xi_t - \langle \mathbf{p}_t, \mathbf{a}_t - \mathbf{d}_t \rangle \mid \mathbf{q}_t\Big] \\ =\ & \mathbb{E}\Big[v(\theta_t, a(\theta_t, \mathbf{q}_t))\xi_t - \sum_{j \in \mathcal{J}} p_{j,t}(a_{j,t} - d_{j,t}) \mid \mathbf{q}_t\Big] \\ =\ & \mathbb{E}\Big[\max_{j \in \mathcal{J}_\emptyset} \left[v(\theta_t, j) - p_{j,t}\right] \xi_t + \sum_{j \in \mathcal{J}} p_{j,t} d_{j,t} \mid \mathbf{q}_t\Big] \\ =\ & \mathbb{E}[\xi_t] \mathbb{E}\Big[\max_{j \in \mathcal{J}_\emptyset} \left[v(\theta_t, j) - p_{j,t}\right] \mid \mathbf{q}_t\Big] + \mathbb{E}\Big[\sum_{j \in \mathcal{J}} p_{j,t} d_{j,t} \mid \mathbf{q}_t\Big] \\ =\ & \frac{\lambda}{1+\lambda} \int_\Theta \max_{j \in \mathcal{J}_\emptyset} \left[v(\theta_t, j) - p_{j,t}\right] dF(\theta) + \frac{1}{1+\lambda} \sum_{j \in \mathcal{J}} \mu_j p_{j,t} \\ =\ & \frac{\lambda}{1+\lambda} h(\mathbf{p}_t) . \end{aligned} \tag{32}$$

Adding $v(\theta_t, a(\theta_t, \mathbf{q}_t))\xi_t$ to both sides of equation (31), we have

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

Taking expectation conditional on $\mathbf{q}_t$ and applying equation (32), we have

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

Rearranging the terms, and we obtain the desired inequality. □

Lemma 10. *For the model defined in Section 3, we have that for any* $\mathbf{q}_t$*,*

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

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 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 definition of Bregman divergence, we have

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

Here the second last inequality follows from the strong smoothness of $L(\mathbf{q})$. This concludes the proof. □

C Price Change Granularity of Nonlinear Waiting Cost Functions

We stated our main result (Theorem 1) in terms of the price change granularity Δ. However, for nonlinear waiting costs, it remains to be shown how Δ is related to the waiting cost function $c(w)$. In this section, we focus on the waiting costs that satisfy the assumption below.

Assumption 2. *We consider the following classes of waiting cost functions.*

- Convex waiting costs. *$c(w)$ is convex, twice-differentiable for $w \geq 0$, and that $c'(w)$ 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 7. *Consider the asymptotic regime in Corollary 2 and waiting cost functions satisfying Assumption 2. 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 7. 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

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

Compare 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 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 non-decreasing, hence

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

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 2, for any $\alpha \in (1, \infty)$ we can apply Lemma 11 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

$$\Delta \leq \frac{2c'\left(\frac{q_{\max,\ell}}{\ell \mu_j}\right)}{\ell \mu_j} \leq \frac{2c'\left(c^{-1}(v_{\max})\right)}{\ell \mu_j}. \tag{33}$$

For concave $c(w)$, using its concavity we have for any $0 \leq q_j \leq q_{\max,\ell}$,

$$p_j(q_j + 1) - p_j(q_j) = \mathbb{E}[c(S_{q_j} + X_{q_j+1}) - c(S_{q_j})] \leq \mathbb{E}[c'(0)X_1] = \frac{c'(0)}{\ell \mu_j}. \tag{34}$$

Combine (33) and (34), we conclude the proof. □

Lemma 11. *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 $\alpha \in (0, \infty)$. For any continuously differentiable function $f(x)$ defined on $\mathbb{R}^+$ such that there exists $C_1, C_2 \in (0, \infty)$ and $f(x) \leq C_1 e^{C_2 x}$, $f'(x) \leq C_1 e^{C_2 x}$ 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 11. The result simply follows the proof of Theorem 1(c) in Kozakiewicz et al. (1947), therefore we omit the details. □

D Proof of Exponentially Small Loss

In this section, we prove the results in Section 5.

D.1 Rate of price adjustment

Proof of Lemma 7. 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 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

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

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 a type θ agent. 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, hence she will choose an item in $\mathcal{J}_\theta^*$ that is the cheapest only if the price is lower than the optimal price, otherwise she will not choose any item. $\mathbf{x}'_\theta$ 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 $-\frac{1}{\lambda}\mathbf{s}' \in \partial h(\mathbf{p})$, 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 agent of type $\theta \in \Theta^*$, under price $\mathbf{p}$ she must choose an item from $\mathcal{J}_\theta^*$, hence

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

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

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

Thus we have

$$\begin{aligned}
\lambda(h(\mathbf{p}) - h(\mathbf{p}^*)) = & \sum_{\theta \in \Theta^*} \lambda_\theta \left(\max_{\mathbf{x}_\theta \in \Delta_\theta} \sum_{j \in \mathcal{J}} (p_j^* - p_j) x_{\theta,j} \right) + \sum_{\theta \in \Theta \setminus \Theta^*} \lambda_\theta \left(\max_{\mathbf{x}_\theta \in \tilde{\Delta}_\theta} \sum_{j \in \mathcal{J}} (p_j^* - p_j) x_{\theta,j} \right) \\
& + \sum_{j \in \mathcal{J}} \mu_j (p_j - p_j^*) \,. \qquad (35)
\end{aligned}$$

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

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

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 (35) 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 concludes 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 (35), we can show that the minimizer of $h(\mathbf{p})$ is non-unique, leading to contradiction with Assumption 1. Note that this already leads to a lower bound of $h(\mathbf{p}) - h(\mathbf{p}^*)$: we have $\epsilon \frac{\mathbf{p}^* - \mathbf{p}}{||\mathbf{p}^* - \mathbf{p}||_2} \subset \mathcal{S}$, 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 that

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

Here the asterisk outside stands for polar set. Denote the LHS set as $\mathcal{B}^*$, 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 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

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

□

D.2 Proof of Theorem 2

Building on the observations obtained in the last two sections, we first prove a lemma that establishes the exponential concentration of $\mathbf{p}$ around $\mathbf{p}^*$, then complete the proof of Theorem 2.

Lemma 12 (Concentration of dual prices)**.** *Suppose Assumption 1 holds. Then for any* $c \leq \frac{\mu_{\min} \delta \gamma}{36}$*, we have*

$$\mathbb{P}(\mathbf{p}_\infty \notin \mathcal{P}) \leq \exp\left(- \log\left(1 + \frac{\gamma \mu_{\min}}{4}\right) \left(\frac{\delta}{12\Delta}\right)\right) .$$

Proof. We prove the result using the Lyapunov functions:

$$\bar{L}(\mathbf{p}) = \frac{1}{2} \sum_{j \in \mathcal{J}} \mu_j (p_j - p_j^*)^2 \qquad \text{and} \qquad V(\mathbf{p}) \triangleq \sqrt{L(\mathbf{p})} \,.$$

Using (12), we have

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

Plugging in the result in Lemma 7, we have

$$\mathbb{E}[\bar{L}(\mathbf{p}_{t+1}) | \mathbf{q}_t] - \bar{L}(\mathbf{p}_t) \leq -c\gamma ||\mathbf{p}_t - \mathbf{p}^*||_2 + c\Delta \,. \tag{36}$$

Use the fact that $f(x) = \sqrt{x}$ is concave for $x \geq 0$ so that for $y > x > 0$, $f(y) - f(x) \leq$

$(y-x)f'(x) = \frac{y-x}{2\sqrt{x}}$, we have

$$V(\mathbf{p}_{t+1}) - V(\mathbf{p}_t) \leq \frac{L(\mathbf{p}_{t+1}) - L(\mathbf{p}_t)}{2V(\mathbf{p}_t)}\,. \tag{37}$$

Take conditional expectation given $\mathbf{q}_t$ on both sides of (37) and plug in (36), we have [PQ: current progress, mind the constants.]

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

Now we use a concentration bound from Bertsimas et al. (2001) to prove the desired result. Let $\mathbf{p}_\infty$ be the steady-state distribution of the prices. For $\mathbf{p}_t$ such that

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

we have

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

As a result, $V(\cdot)$ is a Lyapunov function with exception parameter $\frac{2c}{\mu_{\min}\gamma}$ and negative drift $\frac{c\gamma}{4}$. Note that 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 $r = 0, 1, \cdots$,

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

Note that

$$\left\{\mathbf{p} : V(\mathbf{p}) \leq \frac{\delta}{3}\right\} \subset \left\{\mathbf{p} : ||\mathbf{p} - \mathbf{p}^*||_\infty \leq \frac{\delta}{2}\right\}.$$

As a result, for $c \leq \frac{\mu_{\min}\delta\gamma}{36}$, plugging in $r = \frac{\mu_{\min}\delta}{12c}$, we have

$$\begin{aligned}\mathbb{P}(\mathbf{p}_\infty \notin \mathcal{P}) \leq \mathbb{P}\left(V(\mathbf{p}_\infty) > \frac{\delta}{3}\right) &\leq \left(\frac{1}{1+\frac{\gamma\mu_{\min}}{4}}\right)^{\frac{\mu_{\min}\delta}{12c}} \\ &= \exp\left(-\log\left(1+\frac{\gamma\mu_{\min}}{4}\right)\left(\frac{\delta}{12\Delta}\right)\right).\end{aligned}$$

□