

---

# Regret, stability & fairness in matching markets with bandit learners

---

Sarah H. Cen

Massachusetts Institute of Technology

Devavrat Shah

Massachusetts Institute of Technology

## Abstract

Making an informed decision—for example, when choosing a career or housing—requires knowledge about the available options. Such knowledge is generally acquired through costly trial and error, but this learning process can be disrupted by competition. In this work, we study how competition affects the long-term outcomes of individuals as they learn. We build on a line of work that models this setting as a two-sided matching market with bandit learners. A recent result in this area states that it is impossible to simultaneously guarantee two natural desiderata: stability and low optimal regret for all agents. Resource-allocating platforms can point to this result as a justification for assigning good long-term outcomes to some agents and poor ones to others. We show that this impossibility need not hold true. In particular, by modeling two additional components of competition—namely, costs and transfers—we prove that it is possible to simultaneously guarantee four desiderata: stability, low optimal regret, fairness in the distribution of regret, and high social welfare.

## 1 INTRODUCTION

When individuals compete for resources, it is often assumed that every individual has access to knowledge that allows them to determine which resources they prefer. For example, consider choosing a primary care physician or selecting a career path. Some individuals may have insider knowledge about physicians or careers, e.g., due to previous experiences or advice relayed by friends. Others are left in the dark. Those

in the dark can only obtain knowledge through costly trial and error, which can be further frustrated by the presence of competition. For example, if members of a group are repeatedly denied opportunities in academia due to competition, they may never obtain the requisite knowledge to make an informed decision about pursuing a career in academia.

In this work, we are interested in the interplay between *learning* and *competition*. Specifically, we are motivated by the question: How does competition affect an individual’s ability to make informed decisions and ultimately the individual’s long-term outcomes?

To address this question, we study the *two-sided, competitive matching market*. In this setting, there are two types of agents. For linguistic convenience, we call agents on one side of the market *users* and those on the other *providers*. Each user has a set of preferences over providers, and vice versa. In a matching market, users and providers are paired one-to-one in an assignment is known as the *matching*. Since a user can be paired to at most one provider at a given time, and vice versa, there is *competition* between agents whose preferences overlap. Such markets are ubiquitous and include sellers competing to show advertisements to social media users, students competing for internships, residents competing for housing, and more.

In order to study how knowledge affects the outcomes of individuals in competition, we extend this problem to include a *learning* component, as described next.

**Knowledge in matching.** In the standard centralized matching problem, as posed by Gale and Shapley (1962), every user knows their true preferences over providers *a priori*, and vice versa. Agents then report their preferences to a central platform, which assigns a matching. A matching  $M$  is *stable* if there is no user-provider pair such that they prefer to be matched to one another rather than to their matches under  $M$ . Intuitively, when matchings are stable, agents trust the platform in the sense that there is no incentive for them to defect. While stability is desirable for the platform, it does not necessarily reflect overall performance. Indeed, as studied by Axtell and Kimbrough

---

Proceedings of the 25<sup>th</sup> International Conference on Artificial Intelligence and Statistics (AISTATS) 2022, Valencia, Spain. PMLR: Volume 151. Copyright 2022 by the author(s).

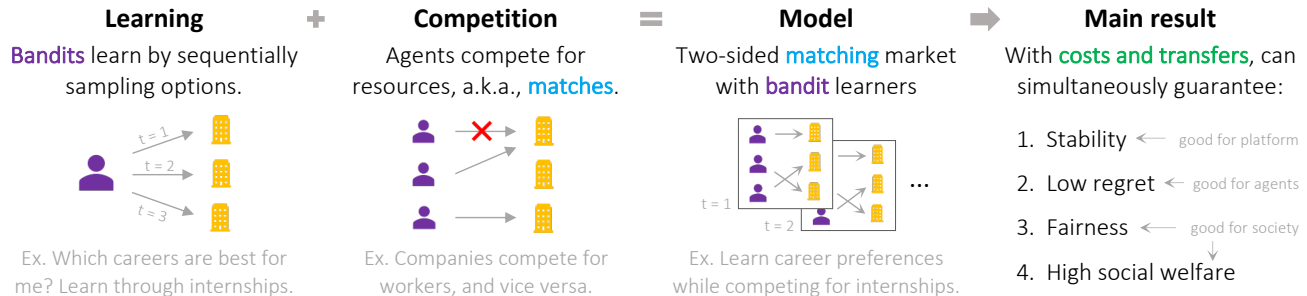


Figure 1: We study learning under competition using the two-sided matching market with bandit learners. As our main result, we show that, with costs and transfers, it is possible to simultaneously guarantee four desiderata.

(2008), stability can come at the expense of other objectives, such as *high social welfare* or *fairness*.

Recent works study an extension of the matching problem in which agents are *not assumed to be knowledgeable*. They build on the insight that, contrary to the core assumption of the standard matching problem, users and providers rarely know their true preferences *a priori*. Rather, the agents learn their preferences through experiences—specifically, by being matched. For example, consider the internship market in which workers present employers with qualifications, and employers present workers with career options. In this setting, matching workers to internships allows workers to learn their preferences over careers and employers to learn their preferences over worker qualifications.

**Blending matching and bandit learning.** While the standard matching problem captures agents competing for resources, a more sophisticated model is needed to model how agents learn through trial and error. Conveniently, the task of *learning one’s preferences from experiences* is at the core of the multi-armed bandit (MAB) problem (Cesa-Bianchi and Lugosi, 2006; Bubeck and Cesa-Bianchi, 2012).

In this way, one can study matching under learning by blending the matching and MAB problems. This is precisely what Liu et al. (2020a) do in a recent work on the centralized two-sided matching problem with bandit learners. Remarkably, Liu et al. (2020a) find that, although agents may not know their true preferences *a priori*, it is possible to assign matchings that are stable at every time step *and* permit agents to learn enough so that the system converges to matchings that are stable under the agents’ *true* (unknown) preferences.

**Our focus.** Liu et al. (2020a) also surface an impossibility result: a matching platform cannot guarantee low *optimal* regret—a measure of long-term happiness—for all agents without sacrificing stability.

This impossibility result is the focus of our work. Consider online platforms (e.g., Facebook), which are gen-

erally concerned with the retention of agents, such as their users, content sources, and advertisers. Recall that a lack of stability implies that agents are better off when they leave the platform. Therefore, online platforms are typically permitted to prioritize stability because it is tied to their financial survival. The impossibility result above gives platforms a justification for favoring some agents over others. In other words, when a platform chooses matches that assign good long-term outcomes to some while assigning poor ones to others, the platform can argue that it could not have done any better by pointing to the impossibility result. *But does this impossibility result hold in general or are there ways to overcome it?*

In this work, we show that the impossibility result is an artifact of the model posed by Liu et al. (2020a), i.e., that stability and low regret need not be in tension. By adding a layer of complexity, we find that it is possible to *simultaneously guarantee stability, low regret, fairness, and high social welfare*.

**Matching with costs and transfers.** Our adjustment is the modeling of costs and transfers. Although this adjustment is minor, it can overturn the impossibility result and align the platform’s and agents’ interests. We provide several examples showing how costs and transfers might arise in real-world settings. Intuitively, competition can give rise to payments—which are equivalent to *transfers* between agents—or demand additional effort and fees from agents—which are equivalent to *costs*.

Our contributions are visualized in Fig. 1 and summarized as follows. We formulate the centralized, two-sided matching problem with bandit learners under costs and transfers in Section 3. In Section 4, we show that, without costs or transfers, it is impossible to guarantee stability alongside low regret, fairness, or high social welfare (Proposition 1). We then show that, with costs and transfers, it is possible to simultaneously guarantee all four desiderata (Theorem 2). In Section 5, we provide intuition via several examples.

## 2 RELATED WORK

**Two-sided matching.** In the two-sided matching problem, there are two distinct populations that wish to be matched one-to-one. Examples of two-sided matchings range from how individuals find partners or how workers fill job openings to how platforms match content to users or how doctors are assigned to patients. One of the challenges is finding a *stable* matching: a matching under which no pair of agents is incentivized to defect and pair off together.

In a foundational work, Gale and Shapley (1962) show that there exists at least one stable matching and provide an algorithm—known as the Gale-Shapley (GS) or deferred acceptance algorithm—that finds such a matching (Knuth and De Bruijn, 1997; Roth, 2008). In the footsteps of Gale and Shapley (1962), one line of research studies *centralized* matchings, in which agents’ preferences are reported to a central platform that assigns the matches. Another line of work studies *decentralized* (or distributed) matching with no coordination between agents (Roth and Vate, 1990; Echenique and Yariv, 2012).

There are many closely related problems, such as the the generalized assignment (Kuhn, 1955; Ross and Soland, 1975; Pentico, 2007); housing allocation (Shapley and Scarf, 1974); college admissions or hospital-residents (Gale and Shapley, 1962; Roth, 1984); stable roommates (Irving, 1985); and generalized stable allocation (Dean and Swar, 2009) problems.

**Multi-armed bandit (MAB) problem.** In the MAB problem (Bubeck and Cesa-Bianchi, 2012; Lattimore and Szepesvári, 2020), an agent (the bandit) sequentially chooses between  $K$  options (or arms). At each time step, the agent selects an arm and receives a noisy reward in return. The agent’s goal is to maximize its cumulative reward over  $T$  time steps or, equivalently, to minimize *regret*. However, the agent does not know the expected rewards of each arm *a priori*. To minimize regret, the bandit learner must balance its short-term desire to select arms with high, known rewards against its need to learn the expected reward of other arms. This balance—also known as the exploitation-exploration trade-off (Wilson et al., 2014; Schnabel et al., 2018; Wilson et al., 2021)—is a focal point of reinforcement learning (Sutton and Barto, 2018; Wiering and Van Otterlo, 2012).

The goal of the MAB problem is to design a selection policy such that the agent incurs low regret. Lai and Robbins (1985) showed that regret must grow  $\Omega(\log(T))$ , which led to research on policies that guarantee  $O(\log(T))$  regret (Thompson, 1933; Bubeck and Cesa-Bianchi, 2012; Lattimore and Szepesvári, 2020).

One of the most popular policies is known as the *upper confidence bound (UCB)* strategy (Auer et al., 2002), which is a deterministic decision-making strategy that achieves low, i.e.,  $O(\log(T))$ , regret.

**Matching markets with bandit learners.** Sequential decision-making under uncertainty and competition is a problem of long-standing interest. Recently, such settings have been modeled as a combined matching and MAB problem (Das and Kamenica, 2005; Liu et al., 2020a). While Das and Kamenica (2005) introduce the combined problem, they study it under the strict assumption that agents on one side of the market have identical preferences. In this work, we study the setting presented by Liu et al. (2020a).

In the centralized matching market with bandit learners (Liu et al., 2020a), users and providers are on opposite sides of the market. As in the matching problem, users compete for matches with providers, and vice versa. As in the MAB problem, agents do not know their true preferences *a priori*. Instead, agents learn their preferences from previous matches. At each time step, each user computes its UCBs for all providers, and vice versa, then reports them to the central platform, which runs the GS algorithm to assign a stable matching at each time step.

There are related works on how competition affects learning (Mansour et al., 2018; Aridor et al., 2020), bandits with collisions (Liu and Zhao, 2010; Kalathil et al., 2014; Bubeck et al., 2020), and coordinated resource allocation (Avner and Mannor, 2019). Though related, these settings are one-sided (i.e., only one side of the market has preferences), and they do not consider costs or transfers. There are also recent works that study other aspects of the matching problem with bandit learners, such as decentralization, matching robustness, and information exchange (Chawla et al., 2020; Boursier and Perchet, 2020; Sankararaman et al., 2020; Vial et al., 2020; Liu et al., 2020b). A recent work by Jagadeesan et al. (2021) study a closely related problem of learning matchings under transferable utilities. These authors consider a different (slightly stronger) notion of stability and they allow transfer rules that vary each time step.

**Costs and transfers.** Recall that, in the MAB problem, an agent’s regret is low if it is  $O(\log(T))$ . Liu et al. (2020a) show that, in the centralized two-sided matching problem with bandit learners, it is possible to guarantee stability and  $O(\log(T))$  *pessimal* regret for all agents. However, there is no guarantee that an agent’s *optimal* regret is  $O(\log(T))$ . We propose this impossibility result can be resolved by adding a layer of complexity not present in the previous work by Liu et al. (2020a): *costs* and *transfers*. Intuitively,

due to competition in matching markets, it is natural for transfers (e.g., payments) to arise or for agents to internalize costs (e.g., fatigue).

Incorporating transfers in matching is not new. Several works on the standard two-sided matching problem allow transfers between agents (Shapley and Shubik, 1971; Becker, 1973). They show that transfers are important to balancing supply and demand, allowing markets to clear (Shapley and Shubik, 1971; Kelso Jr. and Crawford, 1982). It has also been suggested that transfers (e.g., prices) and costs (e.g., search frictions) can explain discrepancies between predicted and observed matching behaviors (Hitsch et al., 2010).

**Stability, fairness, and social welfare.** Early works on matchings study stability (Gale and Shapley, 1962) while later works consider other objectives, such as high social welfare (Axtell and Kimbrough, 2008), fairness (Masarani and Gokturk, 1989; Klaus and Klijn, 2006; Sühr et al., 2019), envy-freeness (Wu and Roth, 2018), and strategy-proofness (Roth, 1982; Ashlagi and Gonczarowski, 2018). While important, stability is a local objective and does not necessarily reflect global performance. For example, stability can come at the expense of social welfare (the sum of all agents’ payoffs) and fairness (the balanced distribution of regret across agents) (Axtell and Kimbrough, 2008). In this work, we consider stability, fairness, and social welfare. We do not study envy-freeness and strategy-proofness. As our main contribution, we show that it is possible to simultaneously guarantee stability, fairness, and high social welfare as well as low regret.

### 3 PROBLEM STATEMENT

In this section, we formalize the centralized matching problem with bandit learners. The setup is similar to that given by Liu et al. (2020a). Our main modification is the incorporation of costs and transfers, and our definitions are adjusted accordingly.

#### 3.1 Setup

Suppose that there are two types of agents who sit on opposite sides of the market. For linguistic convenience, we call the agents on one side *users* and those on the other *providers*. We denote the sets of users and providers by  $\mathcal{U} = \{u_1, u_2, \dots, u_N\}$  and  $\mathcal{P} = \{p_1, p_2, \dots, p_L\}$ , respectively, where  $\mathcal{U} \cap \mathcal{P} = \emptyset$ . Let  $\mathcal{A} = \mathcal{U} \cup \mathcal{P}$  denote the set of all agents and  $\mathcal{A}^+ = \mathcal{A} \cup \{\emptyset\}$ , where  $\emptyset$  denotes no agent. Without loss of generality, let  $N \geq L$ .

Each user has a set of preferences over providers, and vice versa. These *true* preferences are unknown *a priori* and must be learned with time. For example, the

true preference of an applicant for a job could capture how enjoyable the applicant would find the job if given the job, and the true preference of an advertiser for a Facebook user would be how profitable that user is to the advertiser if matched to it. Here, the “agents” are the applicants and jobs in the first example, and the advertisers and Facebook users in the second example.

Let  $\mu(a_1, a_2)$  denote agent  $a_1$ ’s true (unknown) preference for agent  $a_2$ , where  $\mu : \mathcal{A} \times \mathcal{A}^+ \rightarrow \mathbb{R}_{\geq 0}$ , and  $a_1$  and  $a_2$  cannot both belong to  $\mathcal{U}$  or to  $\mathcal{P}$ . We assume that there are no ties, i.e.,  $a' \neq a'' \implies \mu(a, a') \neq \mu(a, a'')$  for all  $a \in \mathcal{A}$ . We also assume that  $\mu(a, a') + \mu(a', a) \neq \mu(a'', a'') + \mu(a'', a''')$  unless  $(a, a') = (a'', a''')$  or  $(a, a') = (a'', a''')$ . This assumption can be relaxed, as we discuss in the Appendix.

**Matching.** Let  $T \in \mathbb{N}_{>0}$  be the time horizon. At every time step  $t \in [T]$ , a platform matches users and providers based on the agents’ preferences  $H_t$  at time  $t$ . Let  $\mathcal{M}(a; H_t)$  denote the agent to which  $a$  is matched at time  $t$ , where  $\mathcal{M} : \mathcal{A} \rightarrow \mathcal{A}^+$  denotes the *matching*, and  $\mathcal{M}(a; H_t) = \emptyset$  implies that agent  $a$  is unmatched at time  $t$ . Without loss of generality, we assume that agents always prefer to be matched than to be unmatched. A *feasible* matching is one in which each user is matched to at most one provider, and all providers are matched to exactly one user. Let  $\mathcal{W}$  be the set of all feasible matchings between users and providers.

An agent  $a$  learns how compatible they are with agent  $a'$  through trial and error. Formally, if  $a$  is matched to  $a'$  at time  $t$ , then  $a$  receives a reward  $X_t(a, a') \stackrel{\text{i.i.d.}}{\sim} \mathcal{D}(a, a')$ , where  $\mathcal{D}(a, a')$  is a sub-Gaussian distribution with parameter  $\sigma^2$  centered at  $\mu(a, a')$ . We adopt the convention that  $\mu(a, \emptyset) = X_t(a, \emptyset) = 0$ . Each agent may also incur a cost and/or partake in a transfer. Let  $\mathcal{C}(a, a'; H_t)$  denote the *cost* that  $a$  would incur if matched to  $a'$  at time  $t$ , where  $\mathcal{C} : \mathcal{A} \times \mathcal{A}^+ \rightarrow \mathbb{R}$ . Let  $\mathcal{T}(a, a'; H_t)$  denote the *transfer* that  $a$  would receive from  $a'$  if matched to  $a'$  at time  $t$ , where  $\mathcal{T} : \mathcal{A} \times \mathcal{A}^+ \rightarrow \mathbb{R}$  and  $\mathcal{T}(a, a'; H_t) = -\mathcal{T}(a', a; H_t)$ . We use the convention that  $\mathcal{C}(a, \emptyset) = \mathcal{T}(a, \emptyset) = 0$  for all  $a \in \mathcal{A}$ .

Therefore, if agent  $a$  is matched to  $a'$  based on preferences  $H_t$ , then  $a$ ’s *observed payoff* is

$$U_t(a, a'; H_t) = X_t(a, a') - \mathcal{C}(a, a'; H_t) + \mathcal{T}(a, a'; H_t).$$

Each agent observes its own match, reward, cost, and transfer but not the private information of any other agent. Other than the matchings, costs, and transfers that it assigns, the platform only observes the preferences  $H_t$  at time  $t$ . The platform cannot observe the agents’ other private information.

**Bandit learners.** We refer to the agents as *bandit learners* because their task—to learn one’s prefer-

erences by sequentially sampling the options—mirrors the MAB problem. Let each agent  $a \in \mathcal{A}$  begin with  $T_0(a, a')$  samples  $\{X^1(a, a'), \dots, X^{T_0(a, a')}(a, a')\}$ , where  $X^\tau(a, a') \stackrel{\text{i.i.d.}}{\sim} \mathcal{D}(a, a')$  for all  $\tau \in [T_0(a, a')]$  and  $T_0(a, a') = T_0(a', a)$ . Let  $T_0(u, p), T_0(p, u) > 0$  for all  $u \in \mathcal{U}$  and  $p \in \mathcal{P}$ . In words, agents have at least one noisy sample before matching begins.<sup>1</sup> Let  $T_t(a, a') = T_0(a, a') + \sum_{\tau \in [t]} \mathbf{1}(\mathcal{M}(a; H_\tau) = a')$  be the number of samples  $a$  has for  $a'$  up to time  $t$  and  $\hat{\mu}_t(a, a') = (T_{t-1}(a, a'))^{-1} (\sum_{\tau \in [t-1]} \mathbf{1}(\mathcal{M}(a; H_\tau) = a') X_\tau(a, a') + \sum_{\tau \in [T_0(a, a')]} X^\tau(a, a'))$  denote the empirical mean of the samples.

Under the  $\sigma^2$ -sub-Gaussian assumption above, the  $\alpha$ -upper confidence bound (UCB) of  $a$  for  $a'$  is:

$$\nu_t(a, a') = \hat{\mu}_t(a, a') + \sqrt{(2\sigma^2\alpha \log(t))/T_{t-1}(a, a')},$$

where  $\alpha \in \mathbb{R} \cap (2, \infty)$  and  $\nu_t(a, \emptyset) = 0$  for all  $a, a' \in \mathcal{A}$ . We call  $\nu_t(a, \cdot)$  the *transient* preferences of agent  $a$  at time  $t$ . To see why, note that in the absence of competition and if  $\mathcal{C} = \mathcal{T} = 0$ , our setup is identical to the classical MAB problem, and  $\nu_t(a, a')$  captures  $a$ 's preference for being matched to (i.e., sampling)  $a'$  at time  $t$  under the UCB algorithm.

In this work, we let  $H_t = \nu_t$ , i.e., agents report their transient preferences  $\nu_t$  to the platform at time  $t$ .<sup>2</sup>

**Summary of matching process.** In summary, the matching process with costs and transfers proceeds as follows. At  $t = 0$ , the platform decides on  $(\mathcal{M}, \mathcal{C}, \mathcal{T})$ , which is made known to all agents. Then, at every time step  $t \in [T]$ , there are four stages:

1. *Update*: Every agent  $a$  updates estimates  $\hat{\mu}_t(a, \cdot)$ .
2. *Report*: Using  $\hat{\mu}_t(a, \cdot)$ , every agent  $a$  reports its transient preferences  $\nu_t(a, \cdot)$  to the platform.
3. *Match and Observe*: Based on  $\nu_t$ , the platform matches users and providers according to  $\mathcal{M}(\cdot; \nu_t)$ . Each agent  $a$  observes its own match  $\mathcal{M}(a; \nu_t)$  and stochastic reward  $X_t(a, \mathcal{M}(a; \nu_t))$ .
4. *Pay and transfer*: Each agent  $a$  may incur the cost  $\mathcal{C}(a, \mathcal{M}(a; \nu_t); \nu_t)$  and/or receive the transfer  $\mathcal{T}(a, \mathcal{M}(a; \nu_t); \nu_t)$  from its match  $\mathcal{M}(a; \nu_t)$ .

As mentioned above, the platform can only observe the agents' transient preferences as well as the matches,

<sup>1</sup>This construction is equivalent (with small modifications) to cyclically matching users and providers in the first  $N$  time steps, as often done in the MAB setting.

<sup>2</sup>Letting  $H_t = \nu_t$  is natural because an agent's preferences at a given time  $t$  must balance exploration and exploitation, as captured by UCB preferences. Note that we could equivalently let  $H_t = \hat{\mu}_t$ , i.e., agents report their empirical estimates of  $\mu$ . Our results would still hold. The only modification would be that the platform computes  $\nu_t$  from  $\hat{\mu}_t$  based on the matching history.

costs, and transfers that it assigns to the agents. An agent can only observe its own information, including its preferences, match, reward, cost, and transfer.

### 3.2 Desiderata

The platform's role is to design the matching mechanism  $\mathcal{M}$ , cost rule  $\mathcal{C}$ , and transfer rule  $\mathcal{T}$ . In this work, there are four desiderata of interest: *stability*, *regret*, *fairness*, and *social welfare*.

We first introduce some notation. We say that  $V(a, a')$  is the payoff that agent  $a$  receives if matched to agent  $a'$  under some payoff function  $V : \mathcal{A} \times \mathcal{A}^+ \rightarrow \mathbb{R}$ . When we are concerned with a specific set of preferences  $\psi : \mathcal{A} \times \mathcal{A}^+ \rightarrow \mathbb{R}$ , we use the following notation:

$$V(a, a'; \psi) = \psi(a, a') - \mathcal{C}(a, a'; \psi) + \mathcal{T}(a, a'; \psi),$$

for all  $a \in \mathcal{A}$  and  $a' \in \mathcal{A}^+$ .

**(a) Stability.** A matching  $M \in \mathcal{W}$  is *not* stable under  $V$  if there exist  $a, a' \in \mathcal{A}$  such that  $V(a, M(a)) < V(a, a')$  and  $V(a', M(a')) < V(a', a)$ .<sup>3</sup> In words, a matching is not stable when there exists a user-provider pair that, under payoffs  $V$ , would rather defect than participate in the matching. Let  $S(V) \subset \mathcal{W}$  be the set of stable matchings under payoffs  $V$ .

**Definition 1.**  $(\mathcal{M}, \mathcal{C}, \mathcal{T})$  is stable if and only if  $\mathcal{M}(\cdot; \nu_t) \in S(V(\cdot, \cdot; \nu_t))$  for all  $t \in [T]$ .

**(b) Low regret.** Regret measures the loss incurred by an agent as it learns. Let the *optimal* and *pessimal* matchings for agent  $a$  at time  $t$  be defined as:

$$\bar{M}_t^a \in \arg \max_{M \in S(V(\cdot, \cdot; \mu))} \mathbb{E}_{X_t \sim \mathcal{D}} [U_t(a, M(a); \nu_t)],$$

$$\underline{M}_t^a \in \arg \min_{M \in S(V(\cdot, \cdot; \mu))} \mathbb{E}_{X_t \sim \mathcal{D}} [U_t(a, M(a); \nu_t)],$$

respectively. For intuition, recall that a stable matching is not necessarily unique, i.e.,  $|S(V(\cdot, \cdot; \mu))| \geq 1$ .  $\bar{M}_t^a$  and  $\underline{M}_t^a$  are both stable matchings under  $V(\cdot, \cdot; \mu)$  that, respectively, yield the highest and lowest expected payoffs for agent  $a$  at time  $t$ . Let the *optimal regret*  $\bar{R}(a; \mathcal{M})$  of agent  $a$  under  $(\mathcal{M}, \mathcal{C}, \mathcal{T})$  be:

$$\sum_{t=1}^T \mathbb{E}[U_t(a, \bar{M}_t^a(a); \nu_t) - U_t(a, \mathcal{M}(a; \nu_t); \nu_t)].$$

Let *pessimal regret* be defined analogously, exchanging  $\bar{M}_t^a$  for  $\underline{M}_t^a$ .<sup>4</sup> Note that  $\bar{R}(a; \mathcal{M}) \geq \underline{R}(a; \mathcal{M})$ . If there is a unique stable matching, i.e.,  $|S(V(\cdot, \cdot; \mu))| = 1$ , then  $\bar{M}_t^a = \underline{M}_t^a$  and  $\bar{R}(a; \mathcal{M}) = \underline{R}(a; \mathcal{M})$  for all  $a \in \mathcal{A}$  and  $t \in [T]$ . We define low regret as follows.

<sup>3</sup>In some works, this definition of stability is known as “weak stability” (Irving, 1994; Manlove, 2002). In this work, when we use “stability”, we mean “weak stability”.

<sup>4</sup>Note that we often suppress  $\mathcal{C}$  and  $\mathcal{T}$  in our notation, such as in  $\bar{R}$ ,  $V$ , and  $W_t$ .

**Algorithm 1:** Gale-Shapley algorithm (with providers as proposers)

**Input:** Set of agents  $\mathcal{A} = \mathcal{U} \cup \mathcal{P}$ , where

$$\begin{aligned} \mathcal{U} &= \{u_1, u_2, \dots, u_N\}, \\ \mathcal{P} &= \{p_1, p_2, \dots, p_L\}, \mathcal{U} \cap \mathcal{P} = \emptyset, \text{ and} \\ N &\geq L. \text{ Payoff function } V : \mathcal{A} \times \mathcal{A}^+ \rightarrow \mathbb{R}. \end{aligned}$$

**Output:** Matching  $M \in S(V) \subset \mathcal{W}$ .

- 1 Initialize  $M : \mathcal{A} \rightarrow \mathcal{A}^+$  s.t.  $M(a) \leftarrow \emptyset \forall a \in \mathcal{A}$ ;
- 2 **while**  $\exists$  unmatched provider  $p \in \mathcal{P} : M(p) = \emptyset$  **do**
- 3     Let  $u$  denote  $p$ 's most preferred user to which  $p$  has not yet proposed;
- 4     If  $u$  is unmatched, i.e.,  $M(u) = \emptyset$ , then match  $u$  and  $p$  such that:  $M(u) \leftarrow p$  and  $M(p) \leftarrow u$ ;
- 5     Otherwise, if  $u$  prefers  $p$  to current match, i.e.,  $V(u, p) > V(u, M(u))$ , match  $u$  and  $p$  s.t.  $M'(M(u)) \leftarrow \emptyset$ ,  $M(u) \leftarrow p$ , and  $M(p) \leftarrow u$ ;
- 6 **end**

The GS algorithm is given in detail in the Appendix.

**Definition 2.** An agent incurs low regret when  $\bar{R}(a; \mathcal{M}) \leq \bar{R}(a; \mathcal{M}) = O(\log(T))$ .

(c) **Fairness.** We say that  $(\mathcal{M}, \mathcal{C}, \mathcal{T})$  is unfair if some agents incur low regret while others incur high regret.

**Definition 3.**  $(\mathcal{M}, \mathcal{C}, \mathcal{T})$  is unfair if there exists a pair of agents  $(a, a')$  such that  $\bar{R}(a; \mathcal{M}) = O(\log(T))$  and  $\bar{R}(a'; \mathcal{M}) = \omega(\log(T))$ . Otherwise, it is fair.

(d) **High social welfare.** Social welfare is a utilitarian metric defined as the sum of all agents' payoffs. Formally, the *social welfare* of  $(\mathcal{M}, \mathcal{C}, \mathcal{T})$  at time step  $t$  is  $W_t(\mathcal{M}) = \sum_{a \in \mathcal{A}} V(a, \mathcal{M}(a; \nu_t); \nu_t)$ .

**Definition 4.**  $(\mathcal{M}, \mathcal{C}, \mathcal{T})$  guarantees  $\kappa$ -high social welfare at time  $t$  if  $W_t(\mathcal{M}) \geq \kappa \max_{\mathcal{M}'} W_t(\mathcal{M}')$  and  $W_t(\mathcal{M}) > 0$ . If there does not exist such a  $\kappa > 0$  or  $W_t(\mathcal{M}) = 0$ , then the social welfare at  $t$  is low.

## 4 MAIN RESULTS

In this section, we give two results. Using our setup, we first re-state the impossibility result of Liu et al. (2020a). We then give our main result, which says that, with costs and transfers, one can simultaneously guarantee stability, low regret, fairness, and high social welfare. All proofs can be found in the Appendix.

Before proceeding, one tool we require is a classical algorithm known as the Gale-Shapley (GS) algorithm, which we restate in Algorithm 1. Several well-known results about the GS algorithm are given in Appendix B. The most relevant result for the understanding of this work is that the GS algorithm always returns a stable matching relative to the given payoff function.

**Impossibility result.** We begin with an impossibility result that is analogous to that given by Liu et al. (2020a) with minor modifications. It states that, in the *absence* of costs and transfers, all agents are guaranteed to have low *peissimal* regret, but there is no guarantee that they have low *optimal* regret, which can grow  $\Omega(T)$ . Let:

$$\begin{aligned} \Delta_{\max}(a) &= \max_{a_1, a_2 \in \mathcal{A}^+} (\mu(a, a_1) - \mu(a, a_2)), \\ \Delta_{\min} &= \min_{a_1 \in \mathcal{A}, a_2 \in \mathcal{A}^+, a_3 \in \mathcal{A}^+ \setminus a_2} |\mu(a_1, a_2) - \mu(a_1, a_3)|. \end{aligned}$$

**Proposition 1.** Suppose that there are no costs or transfers such that:  $\mathcal{C}(a_1, a_2; \psi) = 0$  and  $\mathcal{T}(a_1, a_2; \psi) = 0$  for all  $a_1 \in \mathcal{A}$ ,  $a_2 \in \mathcal{A}^+$ , and  $\psi : \mathcal{A} \times \mathcal{A}^+ \rightarrow \mathbb{R}$ . If the GS algorithm is applied over  $V(\cdot, \cdot; \nu_t)$  at every  $t \in [T]$ , then the system is stable. Moreover,  $\bar{R}(a; \mathcal{M}) \leq 2N^2 L \Delta_{\max}(a) \left( \frac{8\sigma^2 \alpha \log(T)}{\Delta_{\min}^2} + \frac{\alpha}{\alpha-2} \right)$  for all  $a \in \mathcal{A}$ . However, under stability, it is not possible to guarantee fairness or high social welfare, and there exist settings  $(\mathcal{A}, \mu)$  for which  $\bar{R}(a; \mathcal{M}) = \Omega(T)$  for at least one agent  $a \in \mathcal{A}$ .

Proposition 1 states that applying the GS algorithm at each time step ensures that every matching is stable and that every agent has low pessimal regret. However, the second half of Proposition 1 reveals an impossibility result: without costs or transfers, it is not possible to guarantee low *optimal* regret alongside stability. In fact, there are settings in which the optimal regret of at least one agent grows  $\Omega(T)$ . Guaranteeing low pessimal regret reassures pessimistic agents that they will at least reach their worst-case performance under a true stable matching in  $O(\log(T))$  time steps. However, because low optimal regret is not guaranteed, any optimistic agent will be disappointed: it can take  $\Omega(T)$  time steps to surpass their worst-case performance under a true stable matching.

**Incorporating costs and transfers.** Therefore, without costs or transfers, the platform cannot simultaneously guarantee stability and fairness or high social welfare. This negative result begs the question: Is it possible to do better and guarantee low optimal regret, fairness, or high social welfare alongside stability? It turns out that the answer is yes.

**Theorem 2.** There exist cost and transfer rules  $\mathcal{C}(\cdot, \cdot; \nu_t)$  and  $\mathcal{T}(\cdot, \cdot; \nu_t)$  such that, if the GS algorithm is applied over  $V(\cdot, \cdot; \nu_t)$  at every  $t \in [T]$ , then stability, low regret, fairness, and  $\frac{1}{2}$ -high social welfare are guaranteed for all  $(\mathcal{A}, \mu)$ .

Theorem 2 states that, with costs and/or transfers, it is possible to simultaneously guarantee all four desiderata. It implies that the impossibility result of Proposition 1 only holds when costs and transfers are not

allowed. We give intuition for the proof in Section 6. Several questions arise, such as, what types of costs and transfers provide guarantees on stability and regret? Moreover, do such costs and transfers arise in practice? In the following section, we provide insights into these questions by studying a few examples.

## 5 EXAMPLE SETTINGS

In the previous section, we found that, without costs and transfers, it is impossible to simultaneously guarantee stability and low optimal regret. On the other hand, with costs and transfers, it becomes possible to achieve all four desiderata described in Section 3.2. In this section, we unpack several examples in order to understand when this outcome is possible as well as give intuition for when costs and transfers arise. As before, all proofs can be found in the Appendix.

We emphasize that our objective is not to issue recommendations on how to implement costs and transfers. Rather, our goal is to highlight the role of costs and transfers in competitive settings. Intuitively, transfers can be viewed as payments between agents (e.g., prices or distance traveled), and costs can be viewed as frictions (e.g., time expended or entry fees). Neither are necessarily monetary, and both naturally surface in competitive settings. In the following examples, we study three of many possible cost and transfer rules.

### 5.1 Proportional cost with no transfer

In this section, we study cost and transfer rules and show that, while appealing to the platform, they fail to guarantee that low regret holds alongside stability. In the proportional-cost, no-transfer setting,

$$\mathcal{C}(a_1, a_2; \psi) = \gamma\psi(a_1, a_2) \quad \mathcal{T}(a_1, a_2; \psi) = 0, \quad (1)$$

for all  $a_1, a_2 \in \mathcal{A}$  and  $\psi : \mathcal{A} \times \mathcal{A}^+ \rightarrow \mathbb{R}$ , where  $\gamma \in [0, 1]$ . Intuitively, the amount that an agent  $a_1$  expends is proportional to its desire to be matched to  $a_2$ . For example, the amount that a retailer pays a platform to have its advertisement shown to a user or the amount of effort a worker puts into an internship depends on the respective agent's level of interest. The following result shows that proportional costs do not improve upon the no-cost, no-transfer setting.

**Proposition 3.** *Suppose that  $\mathcal{C}$  and  $\mathcal{T}$  are set according to (1). If the GS algorithm is applied over  $V(\cdot, \cdot; \nu_t)$  at every  $t \in [T]$ , then the system is stable. If  $\gamma \in [0, 1)$ ,  $\underline{R}(a; \mathcal{M}) \leq 2N^2L(1 - \gamma)\Delta_{\max}(a) \left( \frac{8\sigma^2\alpha \log(T)}{(1-\gamma)^2\Delta_{\min}^2} + \frac{\alpha}{\alpha-2} \right)$  and, if  $\gamma = 1$ ,  $\underline{R}(a; \mathcal{M}) \leq 0$ . However, under stability, it is not possible to guarantee fairness or high social welfare,*

*and there exist settings  $(\mathcal{A}, \mu)$  for which  $\bar{R}(a; \mathcal{M}) = \Omega(T)$  for at least one agent  $a \in \mathcal{A}$ .*

Proposition 3 shows that this intuitively pleasing setting unfortunately yields the same negative result as the no-cost, no-transfer setting. Interestingly, this setting is appealing to platforms, as shown next.

**Corollary 4.** *Suppose that  $\mathcal{C}$  and  $\mathcal{T}$  are set according to (1) and that the platform wishes to maximize  $F_t : \mathcal{W} \rightarrow \mathbb{R}$  at every  $t \in [T]$ . Then, when  $\gamma = 1$ , there is no performance cost to stability: i.e.,  $\arg \max_{M \in \mathcal{W}} F_t(M) \in S(V(\cdot, \cdot; \nu_t))$  for all  $t \in [T]$ .*

In other words, when  $\gamma = 1$ , the platform can maximize any objective function it desires while remaining confident that the users and providers will continue to use the platform's services. However, by the impossibility result given in Proposition 3, the fact that users and providers continue to use the platform for matching does *not* mean that they do so happily. Some agents may be unhappy because the platform's matchings cause them to incur high optimal regret. Furthermore, they may incur high regret while other agents are allowed to incur low regret. This finding suggests that such a platform may be barely holding onto its agents and may lose them if an alternate platform that offers slightly better outcomes arises.

This setting also illustrates why it is prudent to consider the agents' optimal regret in addition to their pessimal regret. By Proposition 3, an agent's pessimal regret can be non-positive while its optimal regret grows linearly in  $T$ . This finding reinforces the notion that it is not enough to upper bound the pessimal regret because there may be a very large gap between optimal and pessimal regrets.

### 5.2 Balanced transfer

We now introduce a transfer rule under which the platform can simultaneously guarantee all four desiderata. In the no-cost and balanced-transfer setting,

$$\mathcal{C}(a_1, a_2; \psi) = 0, \quad (2)$$

$$\mathcal{T}(a_1, a_2; \psi) = (\psi(a_2, a_1) - \psi(a_1, a_2))/2, \quad (3)$$

for all  $a_1, a_2 \in \mathcal{A}$  and  $\psi : \mathcal{A} \times \mathcal{A}^+ \rightarrow \mathbb{R}$ . To gain an intuition for this setting, suppose  $\psi(a_1, a_2) > \psi(a_2, a_1)$  such that agent  $a_1$  benefits more from the match  $(a_1, a_2)$  than  $a_2$  does. Under (2),  $a_1$  compensates for this imbalance by investing resources to make the experience more amenable for  $a_2$ . For example, suppose that an athlete  $a_1$  is more interested in lessons with coach  $a_2$  than  $a_2$  is interested in coaching  $a_1$ . Then,  $a_1$  can offset this difference by offering to pay more for lessons with  $a_2$ . In the literature, the balanced transfer is viewed as compensation or a bargaining solution.

The next result shows that the platform can simultaneously guarantee all four desiderata. Let

$$\rho(a, a') = \frac{1}{2}(\mu(a, a') + \mu(a', a)),$$

$$\Delta_{\min}^{\rho} = \min_{a_1 \in \mathcal{A}, a_2 \in \mathcal{A}^+, a_3 \in \mathcal{A}^+ \setminus a_2} |\rho(a_1, a_2) - \rho(a_1, a_3)|.$$

**Theorem 5.** *Suppose that  $\mathcal{C}$  and  $\mathcal{T}$  are set according to (2)-(3). Then,  $|S(\rho)| = 1$ . Let  $M^{*,\rho}$  be the only element in  $S(\rho)$  and  $\Delta_{\max}^{*,\rho}(a) = \max_{a' \in \mathcal{A}^+} (\rho(a, M^{*,\rho}(a)) - \rho(a, a'))$ . If the GS algorithm is applied over  $V(\cdot, \cdot; \nu_t)$  at every  $t \in [T]$ , then the system is stable, the social welfare is  $\frac{1}{2}$ -high at all  $t \in [T]$ ,  $\mathcal{M}$  is fair, and  $\underline{R}(a; \mathcal{M}) = \bar{R}(a; \mathcal{M}) \leq \Delta_{\max}^{*,\rho}(a) N^2 L \left( \frac{8\sigma^2 \alpha \log(T)}{(\Delta_{\min}^{\rho})^2} + \frac{\alpha}{\alpha-2} \right)$  for all  $a \in \mathcal{A}$ .*

Therefore, under balanced transfers, applying the GS algorithm at every time step simultaneously guarantees stability, low regret, fairness, and high social welfare. Put differently, under balanced transfers between matched agents, if one guarantees that the matchings are stable, one gets low optimal regret, fairness, and high social welfare for free. One would not necessarily expect such a result to hold true when agents compete while learning (or, equivalently, learn while competing). For example, a college student hoping to learn their ideal career path through internships may be hurt by the presence of competition. It is not clear that an assignment could be stable and allow all agents to learn their preferences efficiently, let alone guarantee fairness and high social welfare.

### 5.3 Pricing

In the previous section, we study a transfer rule in which the payment between agents depends on the preferences of both agents. While such transfers are often feasible, there are settings in which the amount transferred cannot be tailored to the parties involved in the match. For example, suppose each provider is selling a product for which the price does not depend on which user purchases it. To capture such contexts, we turn our attention to the pricing setting, in which:

$$\mathcal{C}(p, u; \psi) = c_1 \quad \mathcal{T}(p, u; \psi) = g(p; \psi) \quad (4)$$

$$\mathcal{C}(u, p; \psi) = c_2 \quad \mathcal{T}(u, p; \psi) = -g(p; \psi) \quad (5)$$

for all  $u \in \mathcal{U}$ ,  $p \in \mathcal{P}$ , and  $\psi : \mathcal{A} \times \mathcal{A}^+ \rightarrow \mathbb{R}$ , where  $c_1, c_2 \in \mathbb{R}$  and  $g : \mathcal{P} \rightarrow \mathbb{R}$ .<sup>5</sup> The transfer  $g(p; \psi)$  can be viewed as the “price” of being matched with  $p$  that

<sup>5</sup>Note that the presence of costs  $c_1$  and  $c_2$  is a technical detail. For those interested, recall that, without loss of generality, we assume all agents prefer to be matched than unmatched, and  $\mathcal{C}(a, \emptyset) = \mathcal{T}(a, \emptyset) = 0$  for all  $a \in \mathcal{A}$ .  $c_1$  and  $c_2$  are normalizing constants to ensure that this assumption

is constant across all agents wishing to “purchase” a match with  $p$ . This setting differs from previous ones in that  $\mathcal{C}$  and  $\mathcal{T}$  do not depend on the preferences of both matched agents. Rather, based on  $\nu_t$ ,  $g(p; \nu_t)$  is assigned to providers in a way that is reminiscent of assigning prices to goods.

For example, suppose  $\mathcal{U}$  is a set of patients, each hoping to book an appointment with a doctor in  $\mathcal{P}$ . Suppose the price of an appointment with  $p$  is fixed such that  $g(p; \nu_t)$  does not depend on the specific patient  $u$ . Since there are a limited number of available appointments,  $g(p; \nu_t)$  may depend on the current demand  $\nu_t$ , which can reflect, for instance, different patient needs and different doctor specialties. If an appointment is booked (i.e., a match is made), then the patient pays the doctor (i.e., makes a transfer). With time, patients learn their preferences over doctors, and doctors learn which patients require their services. In such settings, fairness matters greatly, and preferences must be learned from a small number of sparse interactions.

We now show that, under mild conditions, there always exists at least one pricing rule that simultaneously guarantees stability, low regret, and fairness.

**Theorem 6.** *Suppose that  $\mathcal{C}$  and  $\mathcal{T}$  are set according to (4)-(5). If the GS algorithm is applied over  $V(\cdot, \cdot; \nu_t)$  at every  $t \in [T]$ , then the system is stable. Moreover, if there exists a  $B \in \mathbb{R}_{>0}$  such that  $|\mu(u, \cdot)| \leq B$  for all  $u \in \mathcal{U}$ , then  $|S(V(\cdot, \cdot; \mu))| = 1$  and there exist constants  $c_1$  and  $c_2$  as well as a pricing rule  $g$  such that  $\mathcal{M}$  is fair and  $\underline{R}(a; \mathcal{M}) = \bar{R}(a; \mathcal{M}) \leq 2\Delta_{\max}^{*,B}(a) N^2 L \left( \frac{8\sigma^2 \alpha \log(T)}{(\Delta_{\min})^2} + \frac{\alpha}{\alpha-2} \right)$  for all  $a \in \mathcal{A}$ , where  $M^{*,B}$  is the only element in  $S(V(\cdot, \cdot; \mu))$  and  $\Delta_{\max}^{*,B}(a) = 2B(L-1)\mathbf{1}(a \in \mathcal{U}) + \max_{a' \in \mathcal{A}^+} (\mu(a, M^{*,B}(a)) - \mu(a, a'))$ .*

This result shows that, under mild conditions, there is always at least one way to price matches such that the platform simultaneously guarantees stability, low optimal regret, and fairness. Put differently, one can always find a pricing rule such that, if the matching is stable, low optimal regret and fairness come for free. Under the pricing rule with which we show existence in the proof, high social welfare is not guaranteed, but it may be possible under other pricing rules.

Finally, note that the results in this section hold with small modifications, such as if users and providers are switched; if  $\mathcal{C}(u, p; \psi) = g(p; \psi) - c_1$ ,  $\mathcal{C}(p, u; \psi) = 0$ , and  $\mathcal{T} = 0$ ; or if  $\mathcal{C}(u, p; \psi) = c_2 + (1 - \gamma')g(p; \psi)$ ,  $\mathcal{T}(u, p; \psi) = -\gamma'g(p; \psi)$ ,  $\mathcal{C}(p, u; \psi) = c_1 - (1 - \gamma')g(p; \psi)$ , and  $c\mathcal{T}(p, u; \psi) = \gamma'g(p; \psi)$ .

is consistent with our analysis. For instance, in one setting that we examine in the Appendix,  $c_1 = 0$  and  $c_2 < 0$ .  $-c_2$  can be viewed as a baseline utility that all agents gain if they are matched versus when they are unmatched.



## 6 DISCUSSION

**Summary.** In this work, we model matching under learning by blending the matching and MAB problems. As in the two-sided matching problem, users compete for matches with providers, and vice versa. As in the MAB problem, agents do not know their preferences *a priori* and must learn them from experience.

We focus our attention on the *centralized matching market with bandit learners*, extending the results of Liu et al. (2020a). In particular, the framework offered by Liu et al. (2020a) suggests that it is not possible for competitive matching markets to simultaneously guarantee stability and low optimal regret (a measure of long-term happiness). This impossibility result could have harmful consequences, such as providing a justification for platforms to treat some of its agents well while treating others poorly.

In this work, we show that this impossibility does not always hold. We show that, if we account for two natural components of competition—costs and transfers—it is possible to simultaneously guarantee four desiderata: (1) *stability* such that no pair of agents is incentivized to defect from the matching; (2) *low regret* such that all agents incur  $O(\log(T))$  optimal and pessimal regret; (3) the *fair* distribution of regret across agents; and (4) high *social welfare*.

**Proof intuition.** Ensuring  $(\mathcal{M}, \mathcal{C}, \mathcal{T})$  satisfies the first three desiderata consists of three main ingredients. First, the GS algorithm is applied over  $V(\cdot, \cdot; \nu_t)$  at every time step  $t \in [T]$  to guarantee stability. Second,  $\mathcal{C}$  and  $\mathcal{T}$  are chosen such that, under the appropriate conditions,  $|S(V(\cdot, \cdot; \psi))| = 1$  for any  $\psi$ . In other words,  $\mathcal{C}$  and  $\mathcal{T}$  yield a *unique* stable matching given preferences  $\psi$ . Recalling the definition of regret in Section 3.2, uniqueness ensures that it is not possible for the following to occur: the system converges to a stable matching  $M \in S(V(\cdot, \cdot; \mu))$  for which one agent incurs  $O(\log(T))$  regret at the expense of another agent  $a$  incurring  $\Omega(T)$  regret because  $a$ ’s optimal stable matching is  $M' \in S(V(\cdot, \cdot; \mu)) \setminus M$ . Finally, one must be careful that the rules  $\mathcal{C}$  and  $\mathcal{T}$  that induce a unique stable matching do not interfere with agents’ abilities to learn their preferences  $\mu$ , as discussed next.

**Uniqueness is not enough.** Selecting  $(\mathcal{M}, \mathcal{C}, \mathcal{T})$  such that there is a unique stable matching implies that the user- and provider-optimal stable matchings are identical. It therefore implies that the optimal and pessimal regret are the same. What it does *not* imply, is that the optimal and pessimal regret are  $O(\log(T))$ . One may be tempted to draw the following conclusion: Since Liu et al. (2020) show that pessimal regret is  $O(\log(T))$ , the optimal regret should also be  $O(\log(T))$  when there is a unique stable matching.

However, drawing this conclusion would be a mistake. Why? Because the very tool used to induce a unique stable matching—specifically,  $\mathcal{C}$  and  $\mathcal{T}$ —can interfere efficient learning and prevent the system from converging to a matching that is stable under  $\mu$ . In other words, when there are costs and transfers, even the pessimal regret is not guaranteed to be  $O(\log(T))$ . The key observation here is that the  $O(\log(T))$  bound on pessimal regret given by Liu et al. (2020)—and re-stated in Proposition 1—is proven for a setting in which  $\mathcal{C}$  and  $\mathcal{T}$  do not exist. When  $\mathcal{C}$  and  $\mathcal{T}$  are used to induce a unique stable matching, it takes extra care to ensure that  $\mathcal{C}$  and  $\mathcal{T}$  do not interfere with learning in the process. Furthermore, recall that  $\mu$  is unknown and only  $\nu_t$  is available. Although the goal is to arrive at a matching that is stable under  $\mu$  in  $O(\log(T))$  time steps, one must do so without knowledge of  $\mu$ . Instead,  $\mathcal{C}$  and  $\mathcal{T}$  must be enforced with respect to the time-varying preferences  $\nu_t$ .

**Takeaways for matching platforms.** This work shows that, under costs and transfers, one can simultaneously guarantee four desiderata of interest. We provide one such rule—the balanced transfer—in Section 5.2. We provide another example in Section 5.3 in which the costs and transfers are reminiscent of the pricing of goods. However, there are many more cost and transfer rules that may be of interest.

One of the takeaways of this work is that it is possible to guarantee stability *without* ensuring low regret, as shown in Proposition 1 and Section 5.1. This fact indicates that the continued participation of agents (e.g., social media users) is *not* a sign that agents are satisfied with the platform’s matchings. Indeed, as discussed in Section 5.1, agents can incur high regret even if they do not leave the platform because, by definition, stability implies that agents are not incentivized to defect. It may be to the long-term benefit of a platform to remedy this issue since it leaves room for other platforms that guarantee low regret alongside stability to attract agents away from the platform.

**Future work.** There are many possible paths for future work, such as considering strategy-proofness or allowing information sharing. One could also study related problems with bandit learners under costs and transfers, including decentralized matchings, matchings with ties, the hospital-residents problem (in which multiple users can be matched to the same provider), and more. In addition to the costs and transfers examined in Sections 5, one could characterize the full set of cost and transfer rules under which each desideratum is achievable. Another research direction would be to apply and test the theoretical findings of this work on real-world systems, such as online labor markets or recommender systems.

## Acknowledgements

We would like to thank our reviewers for their time and thoughtful comments. We would also like to thank Meena Jagadeesan, Alexander Wei, Karthik A. Sastry, Zachary Schiffer, and especially Horia Mania for their feedback. This work was supported in parts by the MIT-IBM project on “Representation Learning as a Tool for Causal Discovery”, the NSF TRIPODS Phase II grant towards Foundations of Data Science Institute, the Chyn Duog Shiah Memorial Fellowship, and the Hugh Hampton Young Memorial Fund Fellowship.

## References

- Guy Aridor, Yishay Mansour, Aleksandrs Slivkins, and Zhiwei Steven Wu. Competing Bandits: The Perils of Exploration Under Competition. *arXiv:2007.10144 [cs.GT]*, 2020.
- Itai Ashlagi and Yannai A. Gonczarowski. Stable matching mechanisms are not obviously strategy-proof. *Journal of Economic Theory*, 177:405–425, 2018.
- Peter Auer, Nicolò Cesa-Bianchi, and Paul Fischer. Finite-time Analysis of the Multiarmed Bandit Problem. *Machine learning*, 47(2):235–256, 2002.
- Orly Avner and Shie Mannor. Multi-User Communication Networks: A Coordinated Multi-Armed Bandit Approach. *IEEE/ACM Transactions on Networking*, 27(6):2192–2207, 2019.
- Robert L. Axtell and Steven O. Kimbrough. The High Cost of Stability in Two-Sided Matching: How Much Social Welfare Should be Sacrificed in the Pursuit of Stability. In *Proceedings of the 2008 World Congress on Social Simulation (WCSS 2008)*, Fairfax, USA, 2008.
- Gary S. Becker. A Theory of Marriage: Part I. *Journal of Political Economy*, 81(4):813–846, 1973.
- Etienne Boursier and Vianney Perchet. Selfish Robustness and Equilibria in Multi-Player Bandits. In Jacob Abernethy and Shivani Agarwal, editors, *Proceedings of Thirty Third Annual Conference on Learning Theory (COLT 2020)*, pages 530–581, Online, 2020. PMLR.
- Sébastien Bubeck and Nicolò Cesa-Bianchi. Regret Analysis of Stochastic and Nonstochastic Multi-armed Bandit Problems. *Foundations and Trends in Machine Learning*, 5(1), 2012.
- Sébastien Bubeck, Yuanzhi Li, Yuval Peres, and Mark Sellke. Non-Stochastic Multi-Player Multi-Armed Bandits: Optimal Rate With Collision Information, Sublinear Without. In Jacob Abernethy and Shivani Agarwal, editors, *Proceedings of Thirty Third Annual Conference on Learning Theory (COLT 2020)*, volume 125, pages 961–987, Online, 2020. PMLR.
- Nicolò Cesa-Bianchi and Gábor Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, New York, NY, USA, 2006.
- Ronshee Chawla, Abishek Sankararaman, Ayalvadi Ganesh, and Sanjay Shakkottai. The Gossiping Insert-Eliminate Algorithm for Multi-Agent Bandits. In Silvia Chiappa and Roberto Calandra, editors, *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics (AISTATS 2020)*, pages 3471–3481, Online, 2020. PMLR.
- Sanmay Das and Emir Kamenica. Two-Sided Bandits and the Dating Market. In Leslie Pack Kaelbling, editor, *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence (IJCAI 2005)*, volume 5, pages 947–952, Edinburgh, Scotland, 2005.
- Brian C Dean and Namrata Swar. The Generalized Stable Allocation Problem. In *Proceedings of Third International Workshop on Algorithms and Computation (WALCOM 2009)*, volume 5431, pages 238–249, Kolkata, India, 2009. Springer-Verlag.
- Federico Echenique and Leeat Yariv. An Experimental Study of Decentralized Matching. Technical report, California Institute of Technology, Pasadena, CA, USA, 2012.
- David Gale and Lloyd S. Shapley. College Admissions and the Stability of Marriage. *The American Mathematical Monthly*, 69(1):9–15, 1962.
- Gunter J. Hitsch, Ali Hortaçsu, and Dan Ariely. Matching and Sorting in Online Dating. *The American Economic Review*, 100(1):130–63, 2010.
- Robert W. Irving. An Efficient Algorithm for the “Stable Roommates” Problem. *Journal of Algorithms*, 6(4):577–595, 1985.
- Robert W. Irving. Stable marriage and indifference. *Discrete Applied Mathematics*, 48(3):261–272, 1994.
- Meena Jagadeesan, Alexander Wei, Yixin Wang, Michael I Jordan, and Jacob Steinhardt. Learning equilibria in matching markets from bandit feedback. *arXiv preprint arXiv:2108.08843*, 2021.
- Dileep Kalathil, Naumaan Nayyar, and Rahul Jain. Decentralized Learning for Multiplayer Multiarmed Bandits. *IEEE Transactions on Information Theory*, 60(4):2331–2345, 2014.
- Alexander S. Kelso Jr. and Vincent P. Crawford. Job Matching, Coalition Formation, and Gross Substitutes. *Econometrica*, 50(6):1483–1504, 1982.

- Bettina Klaus and Flip Klijn. Procedurally fair and stable matching. *Economic Theory*, 27(2):431–447, 2006.
- Donald Ervin Knuth and N. G. De Bruijn. *Stable Marriage and Its Relation to Other Combinatorial Problems: An Introduction to the Mathematical Analysis of Algorithms*, volume 10 of *CRM proceedings & lecture notes*. American Mathematical Society, 1997.
- Harold W. Kuhn. The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2(1-2):83–97, 1955.
- T. L. Lai and Herbert Robbins. Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics*, 6(1):4–22, 1985.
- Tor Lattimore and Csaba Szepesvári. *Bandit Algorithms*. Cambridge University Press, Cambridge, UK, 2020.
- Keqin Liu and Qing Zhao. Distributed Learning in Multi-Armed Bandit With Multiple Players. *IEEE Transactions on Signal Processing*, 58(11):5667–5681, 2010.
- Lydia T. Liu, Horia Mania, and Michael I. Jordan. Competing Bandits in Matching Markets. In Silvia Chiappa and Roberto Calandra, editors, *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics (AISTATS 2020)*, pages 1618–1628, Online, 2020a. PMLR.
- Lydia T. Liu, Feng Ruan, Horia Mania, and Michael I. Jordan. Bandit Learning in Decentralized Matching markets. *arXiv:2012.07348 [cs.LG]*, 2020b.
- David F. Manlove. The structure of stable marriage with indifference. *Discrete Applied Mathematics*, 122(1-3):167–181, 2002.
- Yishay Mansour, Aleksandrs Slivkins, and Zhiwei Steven Wu. Competing Bandits: Learning Under Competition. In Anna R. Karlin, editor, *9th Innovations in Theoretical Computer Science Conference (ITCS 2018)*, volume 94 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 48:1–48:27, Dagstuhl, Germany, 2018. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- F. Masarani and Sadik S. Gokturk. On the existence of fair matching algorithms. *Theory and Decision*, 26(3):305–322, 1989.
- David W. Pentico. Assignment problems: A golden anniversary survey. *European Journal of Operational Research*, 176(2):774–793, 2007.
- G. Terry Ross and Richard M. Soland. A branch and bound algorithm for the generalized assignment problem. *Mathematical Programming*, 8:91–103, 1975.
- Alvin E. Roth. The Economics of Matching: Stability and Incentives. *Mathematics of Operations Research*, 7(4):617–628, 1982.
- Alvin E. Roth. The Evolution of the Labor Market for Medical Interns and Residents: A Case Study in Game Theory. 92(6):991–1016, 1984.
- Alvin E Roth. Deferred acceptance algorithms: history, theory, practice, and open questions. *International Journal of Game Theory*, 36:537–569, 2008.
- Alvin E. Roth and John H. Vande Vate. Random Paths to Stability in Two-Sided Matching. *Econometrica*, 58(6):1475–1480, 1990.
- Abishek Sankararaman, Soumya Basu, and Karthik Abinav Sankararaman. Dominate or Delete: Decentralized Competing Bandits with Uniform Valuation. *arXiv:2006.15166 [cs.LG]*, 2020.
- Tobias Schnabel, Paul N. Bennett, Susan T. Dumais, and Thorsten Joachims. Short-Term Satisfaction and Long-Term Coverage: Understanding How Users Tolerate Algorithmic Exploration. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining (WSDM 2018)*, pages 513–521, Marina Del Rey, CA, USA, 2018.
- Lloyd Shapley and Herbert Scarf. On cores and indivisibility. *Journal of Mathematical Economics*, 1: 23–37, 1974.
- Lloyd S. Shapley and Martin Shubik. The Assignment Game I: The Core. *International Journal of Game theory*, 1(1):111–130, 1971.
- Tom Sühr, Asia J. Biega, Meike Zehlike, Krishna P. Gummadi, and Abhijnan Chakraborty. Two-Sided Fairness for Repeated Matchings in Two-Sided Markets: A Case Study of a Ride-Hailing Platform. In *Proceedings of the Twenty Fifth ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 3082–3092, New York, NY, USA, 2019. Association for Computing Machinery.
- Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. The MIT Press, 2nd edition, 2018.
- William R. Thompson. On the Likelihood that One Unknown Probability Exceeds Another in View of the Evidence of Two Samples. *Biometrika*, 25(3-4): 285–294, 1933.
- Roman Vershynin. *High-Dimensional Probability: An Introduction with Applications in Data Science*, volume 47 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, Cambridge, UK, 2018.

- Daniel Vial, Sanjay Shakkottai, and R. Srikant. Robust Multi-Agent Multi-Armed Bandits. *arXiv:2007.03812 [cs.LG]*, 2020.
- Marco Wiering and Martijn Van Otterlo. Reinforcement Learning: State-of-the-Art. *Adaptation, Learning, and Optimization*, 12(3), 2012.
- Robert C. Wilson, Andra Geana, John M. White, and Jonathan D. Ludvig, Elliot A .and Cohen. Humans use directed and random exploration to solve the explore–exploit dilemma. *Journal of Experimental Psychology: General*, 143(6):2074–2081, 2014.
- Robert C. Wilson, Elizabeth Bonawitz, Vincent D. Costa, and R. Becket Ebitz. Balancing exploration and exploitation with information and randomization. *Current Opinion in Behavioral Sciences*, 38: 49–56, 2021.
- Qingyun Wu and Alvin E. Roth. The lattice of envy-free matchings. *Games and Economic Behavior*, 109:201–211, 2018.

---

## Supplementary Material:

# Regret, stability & fairness in matching markets with bandit learners

---

All results that appear in the manuscript are numbered in the same manner below. Results that appear in the Appendix but not in the manuscript are preceded by the letter of the corresponding section.

## A TECHNICAL DETAILS

We re-state the GS algorithm given in Algorithm 1 in detail below. In this algorithm, the providers are the proposers. The version of the algorithm in which users are the proposers is given in Appendix F.

---

**Algorithm 2:** Gale-Shapley algorithm (with providers as proposers)

---

**Input:** Set of agents  $\mathcal{A} = \mathcal{U} \cup \mathcal{P}$ , where the set of users is  $\mathcal{U} = \{u_1, u_2, \dots, u_N\}$ , the set of providers is  $\mathcal{P} = \{p_1, p_2, \dots, p_L\}$ ,  $\mathcal{U} \cap \mathcal{P} = \emptyset$ , and  $N \geq L$ . A payoff function  $V : \mathcal{A} \times \mathcal{A}^+ \rightarrow \mathbb{R}$ .

**Output:** Matching  $M \in S(V) \subset \mathcal{W}$ .

```

1 Initialize matching  $M : \mathcal{A} \rightarrow \mathcal{A}^+$  such that  $M(a) \leftarrow \emptyset$  for all  $a \in \mathcal{A}$ ;
2 Initialize empty (FIFO) queues  $Q(p) \leftarrow []$  for all  $p \in \mathcal{P}$ ;

// Fill each provider's queue with users in order of decreasing preference.
3 for  $p \in \mathcal{P}$  do
4   for  $i = 1, 2, \dots, N$  do
5     Append  $r^{-1}(i; V(p, \cdot))$  to  $Q(p)$ ; // Add  $p$ 's  $i$ -th ranked user.
6   end
7 end

// As long as there exists a provider who is unmatched...
8 while  $\exists p \in \mathcal{P} : M(p) = \emptyset$  do
9    $u \leftarrow \text{pop}(Q(p))$ ; // Provider  $p$ 's favorite user of those remaining in  $p$ 's queue.
// If user  $u$  is unmatched, match  $u$  and  $p$ .
10  if  $M(u) = \emptyset$  then
11     $M(u) \leftarrow p$ ;
12     $M(p) \leftarrow u$ ;

// If user  $u$  prefers  $p$  to its current match  $M(u)$ , match  $u$  and  $p$ .
13  else if  $V(u, p) > V(u, M(u))$  then
14     $M'(M(u)) \leftarrow \emptyset$ ;
15     $M(u) \leftarrow p$ ;
16     $M(p) \leftarrow u$ ;
17 end
18 Return  $M$ ;
```

---

In the GS algorithm presented here, the providers are the proposers. The version of the algorithm in which users are the proposers is given in Appendix F.

**Shorthands.** In this Appendix, we make use of the following notational shorthands. Let  $\psi : \mathcal{A} \times \mathcal{A}^+ \rightarrow \mathbb{R}$  denote an arbitrary function (e.g., preferences). We say that  $a_1 \succ_{\psi(a_3, \cdot)} a_2$  if  $\psi(a_3, a_1) > \psi(a_3, a_2)$ . We denote the rank—or ordinal preference—of agent  $a'$  under  $\psi(a, \cdot)$  by  $r(a'; \psi(a, \cdot))$ , where  $r : \mathcal{A}^+ \rightarrow \mathbb{N}_{>0}$  and  $r(a'; \psi(a, \cdot)) < r(a''; \psi(a, \cdot)) \iff \psi(a, a') > \psi(a, a'')$ . For example, if  $p = \arg \max_{p \in \mathcal{P}} \psi(u, p)$ , then  $r(p; \psi(u, \cdot)) = 1$ . Accordingly,  $r^{-1}(n; \psi(a, \cdot))$  is the  $n$ -th highest ranked user under  $\psi(a, \cdot)$ .

**Assumptions.** At various points in our analysis, we make uniqueness assumptions. These assumptions are fairly mild, as they require that the corresponding real-valued preferences (or payoffs) differ by some amount, however small. Intuitively, an assumption of uniqueness is equivalent to an assumption that an agent has a strict ordering of preferences (or payoffs). This assumption is therefore not valid when there are ties (i.e., indifference) in an agent’s preferences (or payoffs). We can also relax one of our assumptions that  $\mu(a, a') + \mu(a', a) \neq \mu(a''', a'') + \mu(a'', a''')$  unless  $(a, a') = (a'', a''')$  or  $(a, a') = (a'', a''')$ . This assumption is only relevant to Theorem 2 and Theorem 5. Even so, it is not strictly necessary. The assumption could be removed by introducing a consistent tie-breaking rule that the platform uses to choose between matchings, as discussed in Section F. Because this setup requires additional notation, we opt for the more straightforward construction above.

## B PROOFS FOR STABILITY RESULTS

**Lemma B.1.** *Let  $\psi : \mathcal{A} \times \mathcal{A}^+ \rightarrow \mathbb{R}$  denote an arbitrary preference function. Let  $\mathcal{M}(\cdot; \psi)$  denote a matching returned by the GS algorithm, as given in Algorithm 2, when applied over payoffs  $V(\cdot, \cdot; \psi)$ . Then,  $\mathcal{M}(\cdot; \psi)$  is a stable matching under payoffs  $V(\cdot, \cdot; \psi)$ : i.e.,  $\mathcal{M}(\cdot; \psi) \in S(V(\cdot, \cdot; \psi))$ .*

*Proof.* This is a well-known result, cf. Gale and Shapley (1962) or Roth (2008). To prove it, consider Algorithm 2. Suppose that the result is not true, in which case there exists a  $p_1 \in \mathcal{P}$ ,  $p_2 \in \mathcal{P} \cup \emptyset$ , and  $u_1, u_2 \in \mathcal{U}$  such that  $\mathcal{M}(u_1; \psi) = p_1$  and  $\mathcal{M}(u_2; \psi) = p_2$ , but  $u_2 \succ_{V(p_1, \cdot; \psi)} u_1$  and  $p_1 \succ_{V(u_2, \cdot; \psi)} p_2$  by the definition of stability (see Section 3.2).

$u_2 \succ_{V(p_1, \cdot; \psi)} u_1$  implies that  $p_1$  proposes to  $u_2$  before proposing to  $u_1$ . However, since  $p_1 \succ_{V(u_2, \cdot; \psi)} p_2$ , it is not possible that (a)  $u_2$  rejects  $p_1$  in favor of  $p_2$  or (b)  $u_2$  rejects  $p_1$  in favor of  $p_3 \in \mathcal{P} \setminus \{p_1, p_2\}$  but that  $u_2$  does not reject  $p_2$  in favor of  $p_3$ . Therefore, by contradiction, the GS algorithm returns a stable matching.  $\square$

**Lemma B.2.** *Consider the setup in Section 3.1. Suppose that, at every time step  $t \in [T]$ ,  $\mathcal{M}(\cdot; \nu_t)$  is a matching returned by the GS algorithm using the payoffs  $V(\cdot, \cdot; \nu_t)$ . Then,  $(\mathcal{M}, \mathcal{C}, \mathcal{T})$  is stable.*

*Proof.* This result follows from Lemma B.1 and Definition 1.  $\square$

In the remaining analysis, we will also make use of the following definition and lemma.

**Definition B.1.** *A stable matching  $M \in S(V)$  is  $\mathcal{G}$ -optimal under payoffs  $V : \mathcal{A} \times \mathcal{A}^+ \rightarrow \mathbb{R}$  if  $V(a, M(a)) \geq V(a, M'(a))$  for all  $M' \in S(V)$  and  $a \in \mathcal{G}$ . A stable matching  $M \in S(V)$  is  $\mathcal{G}$ -pessimal under payoffs  $V : \mathcal{A} \times \mathcal{A}^+ \rightarrow \mathbb{R}$  if  $V(a, M(a)) \leq V(a, M'(a))$  for all  $M' \in S(V)$  and  $a \in \mathcal{G}$ .*

**Lemma B.3.** *Suppose that the GS algorithm is performed over payoffs  $V : \mathcal{A} \times \mathcal{A}^+ \rightarrow \mathbb{R}$  and with  $\mathcal{G}$  as proposers. Suppose further that  $a' \neq a'' \implies V(a, a') \neq V(a, a'')$  for all  $a \in \mathcal{A}$ . Then, if the matching  $M$  is returned by the GS algorithm,  $M \in S(V)$  is  $\mathcal{G}$ -optimal and  $(\mathcal{A} \setminus \mathcal{G})$ -pessimal.*

*Proof.* This result is well-known, cf. Theorem 3 given by Roth (2008).  $\square$

## C PRELIMINARY RESULTS

Let  $T(\mathcal{E}) = \sum_{t=1}^T \mathbf{1}(\mathcal{E} \text{ occurs at time step } t)$  denote the number of times event  $\mathcal{E}$  occurs over all  $T$  time steps.

**Lemma C.4.** *Consider an agent faced with  $K$  arms, as in the standard multi-armed bandit (MAB) problem. Let  $T \in \mathbb{N}_{>0}$  be the time horizon. Suppose that, at each time step  $t \in [T]$ , the agent is either: (1) matched with arm  $I_t$ , in which case it receives reward  $Y_{I_t, t}$ ; or (2) unmatched, in which case it receives no reward. Suppose  $Y_{i, t}$  is drawn i.i.d. from a sub-Gaussian distribution with parameter  $\sigma^2$  for all  $i \in [K]$ . Let  $m(i) = \mathbb{E}[Y_{i, 0}]$ ,  $m^* = \max_{i \in [K]} m(i)$ , and  $T_t(i) = \sum_{\tau=1}^t \mathbf{1}(I_\tau = i)$ . Let*

$$v_t(i) = \frac{1}{T_{t-1}(i)} \sum_{t=1}^T Y_{i, t} \mathbf{1}(I_t = i) + \sqrt{\frac{2\sigma^2 \alpha \log(t)}{T_{t-1}(i)}},$$

be the agent's upper confidence bound (UCB) for all arms  $i \in [K]$ , where  $\alpha > 2$ . Then,

$$\begin{aligned} \mathbb{E}[T(I_t = i \cap v_t(i) \geq v_t(j) \cap m(i) < m(j))] &= \mathbb{E} \left[ \sum_{t=1}^T \mathbf{1}(I_t = i \cap v_t(i) \geq v_t(j) \cap m(i) < m(j)) \right] \\ &\leq \frac{8\sigma^2\alpha \log(T)}{(m(j) - m(i))^2} + \frac{\alpha}{\alpha - 2}, \end{aligned}$$

for all  $i, j \in [K]$  and  $i \neq j$ , where the expectation is taken with respect to the randomness of rewards  $Y_{I_t, t}$ .

*Proof.* This proof is adapted from the proof that upper bounds regret in the MAB problem under the UCB strategy, cf. Theorem 2.1 presented by Bubeck and Cesa-Bianchi (2012).

Let  $\hat{m}_t(i) = \frac{1}{T_{t-1}(i)} \sum_{\tau=1}^t Y_{i, \tau} \mathbf{1}(I_\tau = i)$ ,  $\Delta(i, j) = |m(i) - m(j)|$ , and  $\Psi^*(\epsilon) := \frac{\epsilon^2}{2\sigma^2}$ . If  $I_t = i$ ,  $v_t(i) \geq v_t(j)$ , and  $m(i) < m(j)$ , then at least one of the following must be true:

$$\hat{m}_t(j) + (\Psi^*)^{-1} \left( \frac{\alpha \log(t)}{T_{t-1}(j)} \right) \leq m(j), \quad (6)$$

$$\hat{m}_t(i) > m(i) + (\Psi^*)^{-1} \left( \frac{\alpha \log(t)}{T_{t-1}(i)} \right), \quad (7)$$

$$T_{t-1}(i) < \frac{\alpha \log(t)}{\Psi^*(\Delta(i, j)/2)}. \quad (8)$$

To confirm that one of the prior statements must be true, suppose that all three statements are false. Then,

$$v_t(j) = \hat{m}_t(j) + (\Psi^*)^{-1} \left( \frac{\alpha \log(t)}{T_{t-1}(j)} \right) \quad (9)$$

$$> m(j) \quad (10)$$

$$= m(i) + \Delta(i, j) \quad (11)$$

$$\geq m(i) + 2(\Psi^*)^{-1} \left( \frac{\alpha \log(t)}{T_{t-1}(i)} \right) \quad (12)$$

$$\geq \hat{m}_t(i) + (\Psi^*)^{-1} \left( \frac{\alpha \log(t)}{T_{t-1}(i)} \right) \quad (13)$$

$$= v_t(i), \quad (14)$$

where (9) and (14) follow from the definition of  $v_t$ ; (10) follows from the assumption that (6) is false; (11) follows from the definition of  $\Delta(i, j)$  and that  $m(j) > m(i)$ ; (12) follows from the assumption that (8) is false; and (13) follows from the assumption that (7) is false. However, (14) implies that  $v_t(j) > v_t(i)$ , which cannot be true by assumption.

Therefore, we have a contradiction, which implies that, if  $I_t = i$ ,  $v_t(i) \geq v_t(j)$ , and  $m(i) < m(j)$ , then at least

one of the statements (6)-(8) must be true. Let  $\bar{T}_{ij} = \lceil \frac{\alpha \log(T)}{\Psi^*(\Delta(i,j)/2)} \rceil$ . Then,

$$\begin{aligned}
 & \mathbb{E}[T(I_t = i \cap v_t(i) \geq v_t(j) \cap m(i) < m(j))] \\
 &= \mathbb{E} \left[ \sum_{t=1}^T \mathbf{1}(I_t = i \cap v_t(i) \geq v_t(j) \cap m(i) < m(j)) \right] \\
 &\leq \mathbb{E} \left[ \sum_{t=1}^T \mathbf{1} \left( I_t = i \cap T_{t-1}(i) < \frac{\alpha \log(t)}{\Psi^*(\Delta(i,j)/2)} \right) \right] \\
 &\quad + \mathbb{E} \left[ \sum_{t=\bar{T}_{ij}+1}^T \mathbf{1} \left( \hat{m}_t(j) + (\Psi^*)^{-1} \left( \frac{\alpha \log(t)}{T_{t-1}(j)} \right) \right. \right. \\
 &\quad \quad \left. \left. \leq m(j) \cup \hat{m}_t(i) > m(i) + (\Psi^*)^{-1} \left( \frac{\alpha \log(t)}{T_{t-1}(i)} \right) \right) \right] \\
 &\leq \mathbb{E} \left[ \sum_{t=1}^T \mathbf{1} \left( I_t = i \cap T_{t-1}(i) < \frac{\alpha \log(T)}{\Psi^*(\Delta(i,j)/2)} \right) \right] \\
 &\quad + \mathbb{E} \left[ \sum_{t=\bar{T}_{ij}+1}^T \mathbf{1} \left( \hat{m}_t(j) + (\Psi^*)^{-1} \left( \frac{\alpha \log(t)}{T_{t-1}(j)} \right) \leq m(j) \right. \right. \\
 &\quad \quad \left. \left. \cup \hat{m}_t(i) > m(i) + (\Psi^*)^{-1} \left( \frac{\alpha \log(t)}{T_{t-1}(i)} \right) \right) \right] \\
 &\leq \bar{T}_{ij} + \mathbb{E} \left[ \sum_{t=\bar{T}_{ij}+1}^T \mathbf{1} \left( \hat{m}_t(j) + (\Psi^*)^{-1} \left( \frac{\alpha \log(t)}{T_{t-1}(j)} \right) \leq m(j) \right. \right. \\
 &\quad \quad \left. \left. \cup \hat{m}_t(i) > m(i) + (\Psi^*)^{-1} \left( \frac{\alpha \log(t)}{T_{t-1}(i)} \right) \right) \right] \\
 &\leq \bar{T}_{ij} + \sum_{t=\bar{T}_{ij}+1}^T \left[ \mathbb{P} \left( \hat{m}_t(j) + (\Psi^*)^{-1} \left( \frac{\alpha \log(t)}{T_{t-1}(j)} \right) \leq m(j) \right) \right. \\
 &\quad \left. + \mathbb{P} \left( \hat{m}_t(i) > m(i) + (\Psi^*)^{-1} \left( \frac{\alpha \log(t)}{T_{t-1}(i)} \right) \right) \right]. \tag{15}
 \end{aligned}$$

It remains to bound the probabilities on the right-hand side in (15). Beginning with the left term under the sum in (15):

$$\begin{aligned}
 & \mathbb{P} \left( \hat{m}_t(j) + (\Psi^*)^{-1} \left( \frac{\alpha \log(t)}{T_{t-1}(j)} \right) \leq m(j) \right) \\
 &\leq \mathbb{P} \left( \exists \tau \in \{1, \dots, t\} : T_{t-1}(j) = \tau \cap \hat{m}_t(j) + (\Psi^*)^{-1} \left( \frac{\alpha \log(t)}{\tau} \right) \leq m(j) \right) \\
 &\leq \sum_{\tau=1}^t \mathbb{P} \left( T_{t-1}(j) = \tau \cap \hat{m}_t(j) + (\Psi^*)^{-1} \left( \frac{\alpha \log(t)}{\tau} \right) \leq m(j) \right) \\
 &\leq \sum_{\tau=1}^t \frac{1}{t^\alpha} \\
 &= \frac{1}{t^{\alpha-1}},
 \end{aligned}$$

where the second-to-last line follows from the concentration properties of sub-Gaussian random variables (Ver-



shynin, 2018). Incorporating the sum in (15),

$$\begin{aligned}
 \sum_{t=\bar{T}_{ij}+1}^T \mathbb{P} \left( \hat{m}_t(j) + (\Psi^*)^{-1} \left( \frac{\alpha \log(t)}{T_{t-1}(j)} \right) \leq m(j) \right) &\leq \sum_{t=\bar{T}_{ij}+1}^T \frac{1}{t^{\alpha-1}} \\
 &\leq \sum_{t=\bar{T}_{ij}+1}^{\infty} \frac{1}{t^{\alpha-1}} \\
 &\leq \sum_{t=2}^{\infty} \frac{1}{t^{\alpha-1}} \\
 &\leq \int_1^{\infty} \frac{1}{t^{\alpha-1}} dt \\
 &= \frac{1}{\alpha-2}.
 \end{aligned} \tag{16}$$

Using a symmetric analysis, the right term under the sum in (15) can be upper bounded by the same quantity, such that

$$\sum_{t=\bar{T}_{ij}+1}^T \mathbb{P} \left( \hat{m}_t(i) > m(i) + (\Psi^*)^{-1} \left( \frac{\alpha \log(t)}{T_{t-1}(i)} \right) \right) \leq \frac{1}{\alpha-2}. \tag{17}$$

Combining with (16) and (17) with (15):

$$\begin{aligned}
 \mathbb{E}[T(I_t = i \cap v_t(i) \geq v_t(j) \cap m(i) < m(j))] &\leq \bar{T}_{ij} + \frac{2}{\alpha-2} \\
 &\leq \frac{\alpha \log(T)}{\Psi^*(\Delta(i, j)/2)} + 1 + \frac{2}{\alpha-2} \\
 &\leq \frac{\alpha \log(T)}{\Psi^*(\Delta(i, j)/2)} + \frac{\alpha}{\alpha-2},
 \end{aligned}$$

which, after substituting in for the definition of  $\Psi^*$ , gives the result as stated in the lemma.

To connect this proof to Theorem 2.1 in Bubeck and Cesa-Bianchi (2012), note that  $\Psi^*$  is the Legendre transform (convex-conjugate) of  $Y - \mathbb{E}[Y]$ , where  $Y$  is a sub-Gaussian random variable with parameter  $\sigma^2$  and that

$$(\Psi^*)^{-1} \left( \frac{\alpha \log(t)}{T_{t-1}(i)} \right) = \sqrt{\frac{2\sigma^2 \alpha \log(t)}{T_{t-1}(i)}},$$

is the additive term in the UCB formula.  $\square$

**Lemma C.5.** *Consider the setup described in Section 3.1. Let  $\psi, \psi' : \mathcal{A} \times \mathcal{A}^+ \rightarrow \mathbb{R}$ . If  $\mathcal{M}(\cdot; \psi') \in S(V(\cdot, \cdot; \psi'))$  and  $\mathcal{M}(\cdot; \psi') \notin S(V(\cdot, \cdot; \psi))$ , then there exists  $a_1, a_2, a_3, a_4 \in \mathcal{A}$  such that  $\mathcal{M}(a_1; \psi') = a_2$  and  $\mathcal{M}(a_3; \psi') = a_4$ , but  $V(a_1, a_4; \psi) > V(a_1, a_2; \psi)$  and  $V(a_4, a_1; \psi) > V(a_4, a_3; \psi)$  and either  $V(a_1, a_4; \psi') \leq V(a_1, a_2; \psi')$  or  $V(a_4, a_1; \psi') \leq V(a_4, a_3; \psi')$ .*

*Proof.* If  $\mathcal{M}(\cdot; \psi') \notin S(V(\cdot, \cdot; \psi))$ , then there exists  $a_1, a_2 \in \mathcal{A}$  such that  $\mathcal{M}(a_1; \psi') = a_2$  but  $(a_1, a_2)$  is not stable under payoffs  $V(\cdot, \cdot; \psi)$ . The latter implies that there exists  $a_3, a_4 \in \mathcal{A}$  such that  $\mathcal{M}(a_3; \psi') = a_4$  and  $V(a_1, a_4; \psi) > V(a_1, a_2; \psi)$  and  $V(a_4, a_1; \psi) > V(a_4, a_3; \psi)$ . However, because the matches  $(a_1, a_2)$  and  $(a_3, a_4)$  are stable under  $V(\cdot, \cdot; \psi')$ , either  $V(a_1, a_4; \psi') \leq V(a_1, a_2; \psi')$  or  $V(a_4, a_1; \psi') \leq V(a_4, a_3; \psi')$ .  $\square$

**Lemma C.6.** *Consider the setup described in Section 3.1. Let  $\Delta_{\min} = \min_{a_1 \in \mathcal{A}, a_2 \in \mathcal{A}^+, a_3 \in \mathcal{A}^+ \setminus a_2} |\mu(a_1, a_2) - \mu(a_1, a_3)|$ . Suppose there is no cost and no transfer such that  $\mathcal{C} = \mathcal{T} = 0$ . Then,*

$$\mathbb{E}T(\mathcal{M}(\cdot; \nu_t) \notin S(V(\cdot, \cdot; \mu))) = \mathbb{E}T(\mathcal{M}(\cdot; \nu_t) \notin S(\mu)) \leq 2N^2L \left( \frac{8\sigma^2 \alpha \log(T)}{\Delta_{\min}^2} + \frac{\alpha}{\alpha-2} \right).$$

*Proof.* Under no cost and no transfer,  $V(\cdot, \cdot; \mu) = \mu(\cdot, \cdot)$  and  $V(\cdot, \cdot; \nu_t) = \nu_t(\cdot, \cdot)$ . Recall that, by construction,  $\mathcal{M}(\cdot; \nu_t) \in S(V(\cdot, \cdot; \nu_t)) = S(\nu_t)$ . If, additionally,  $\mathcal{M}(\cdot; \nu_t) \notin S(V(\cdot, \cdot; \mu)) = S(\mu)$ , then, by Lemma C.5, there exists  $a_1, a_2, a_3, a_4 \in \mathcal{A}$  such that  $\mathcal{M}(a_1; \nu_t) = a_2$ ,  $\mathcal{M}(a_3; \nu_t) = a_4$ ,  $\mu(a_1, a_4) > \mu(a_1, a_2)$ , and  $\mu(a_4, a_1) > \mu(a_4, a_3)$  and either  $\nu_t(a_1, a_4) \leq \nu_t(a_1, a_2)$  or  $\nu_t(a_4, a_1) \leq \nu_t(a_4, a_3)$ . As such,

$$\begin{aligned}
 & \mathbb{E}T(\mathcal{M}(\cdot; \nu_t) \notin S(\mu)) \\
 & \leq \mathbb{E}T(\exists a_1, a_2, a_3, a_4 \in \mathcal{A} : \mathcal{M}(a_1; \nu_t) = a_2 \cap \mathcal{M}(a_3; \nu_t) = a_4 \\
 & \quad \cap \mu(a_1, a_4) > \mu(a_1, a_2) \\
 & \quad \cap \mu(a_4, a_1) > \mu(a_4, a_3) \\
 & \quad \cap (\nu_t(a_1, a_4) \leq \nu_t(a_1, a_2) \cup \nu_t(a_4, a_1) \leq \nu_t(a_4, a_3))) \\
 & \leq \mathbb{E}T(\exists a_1, a_2, a_3, a_4 \in \mathcal{A} : \mathcal{M}(a_1; \nu_t) = a_2 \cap \mathcal{M}(a_3; \nu_t) = a_4 \\
 & \quad \cap \mu(a_1, a_4) > \mu(a_1, a_2) \\
 & \quad \cap \mu(a_4, a_1) > \mu(a_4, a_3) \\
 & \quad \cap \nu_t(a_1, a_4) \leq \nu_t(a_1, a_2)) \\
 & + \mathbb{E}T(\exists a_1, a_2, a_3, a_4 \in \mathcal{A} : \mathcal{M}(a_1; \nu_t) = a_2 \cap \mathcal{M}(a_3; \nu_t) = a_4 \\
 & \quad \cap \mu(a_1, a_4) > \mu(a_1, a_2) \\
 & \quad \cap \mu(a_4, a_1) > \mu(a_4, a_3) \\
 & \quad \cap \nu_t(a_4, a_1) \leq \nu_t(a_4, a_3)) \\
 & \leq \mathbb{E}T(\exists a_1, a_2, a_4 : \mathcal{M}(a_1; \nu_t) = a_2 \cap \mu(a_1, a_4) > \mu(a_1, a_2) \cap \nu_t(a_1, a_4) \leq \nu_t(a_1, a_2)) \\
 & \quad + \mathbb{E}T(\exists a_1, a_3, a_4 : \mathcal{M}(a_3; \nu_t) = a_4 \cap \mu(a_4, a_1) > \mu(a_4, a_3) \cap \nu_t(a_4, a_1) \leq \nu_t(a_4, a_3)) \\
 & = 2\mathbb{E}T(\exists a_1, a_2, a_3 : \mathcal{M}(a_1; \nu_t) = a_2 \cap \mu(a_1, a_3) > \mu(a_1, a_2) \cap \nu_t(a_1, a_3) \leq \nu_t(a_1, a_2)).
 \end{aligned}$$

Then,

$$\begin{aligned}
 & \mathbb{E}T(\mathcal{M}(\cdot; \nu_t) \notin S(\mu)) \\
 & \leq 2\mathbb{E}T(\exists a_1, a_2, a_3 : \mathcal{M}(a_1; \nu_t) = a_2 \cap \mu(a_1, a_3) > \mu(a_1, a_2) \cap \nu_t(a_1, a_3) \leq \nu_t(a_1, a_2)) \\
 & \leq 2 \sum_{a_1, a_2, a_3 \in \mathcal{A}} \mathbb{E}T(\mathcal{M}(a_1; \nu_t) = a_2 \cap \mu(a_1, a_3) > \mu(a_1, a_2) \cap \nu_t(a_1, a_3) \leq \nu_t(a_1, a_2)) \\
 & \leq 2N^2L \left( \frac{8\sigma^2\alpha \log(T)}{\Delta_{\min}^2} + \frac{\alpha}{\alpha-2} \right),
 \end{aligned}$$

where the last line follows from Lemma C.4 since, from the perspective of agent  $a_1$ , being matched with  $a_2$  is identical to  $a_1$  sampling arm  $a_2$  in the MAB problem.  $\square$

## D PROOFS OF MAIN RESULTS (SECTION 4)

Recall that

$$\begin{aligned}
 \Delta_{\max}(a) &= \max_{a_1, a_2 \in \mathcal{A}^+} (\mu(a, a_1), -\mu(a, a_2)) \\
 \Delta_{\min} &= \min_{a_1 \in \mathcal{A}, a_2 \in \mathcal{A}^+, a_3 \in \mathcal{A}^+ \setminus a_2} |\mu(a_1, a_2) - \mu(a_1, a_3)|.
 \end{aligned}$$

**Proposition 1.** *Suppose that there are no costs or transfers such that:  $\mathcal{C}(a_1, a_2; \psi) = 0$  and  $\mathcal{T}(a_1, a_2; \psi) = 0$  for all  $a_1 \in \mathcal{A}$ ,  $a_2 \in \mathcal{A}^+$ , and  $\psi : \mathcal{A} \times \mathcal{A}^+ \rightarrow \mathbb{R}$ . If the GS algorithm is applied over  $V(\cdot, \cdot; \nu_t)$  at every  $t \in [T]$ , then the system is stable. Moreover,  $\underline{R}(a; \mathcal{M}) \leq 2N^2L\Delta_{\max}(a) \left( \frac{8\sigma^2\alpha \log(T)}{\Delta_{\min}^2} + \frac{\alpha}{\alpha-2} \right)$  for all  $a \in \mathcal{A}$ . However, under stability, it is not possible to guarantee fairness or high social welfare, and there exist settings  $(\mathcal{A}, \mu)$  for which  $\bar{R}(a; \mathcal{M}) = \Omega(T)$  for at least one agent  $a \in \mathcal{A}$ .*

*Proof.* By Lemma B.2, the system is stable since the GS algorithm is applied at each time step. Next, we perform the regret analysis.

**Logarithmic pessimal regret.** Recall that  $\underline{M}_t^a \in \arg \min_{M \in S(V(\cdot, \cdot; \mu))} \mathbb{E}[U_t(a, M(a); \nu_t)]$ . Under no cost or transfer,  $\mathbb{E}[U_t(a, a'; \nu_t)] = \mu(a, a')$  and  $V(\cdot, \cdot; \mu) = \mu(\cdot, \cdot)$ . As such,  $\underline{M}_t^a \in \arg \min_{M \in S(\mu)} \mu(a, M(a))$ , and the pessimal regret is

$$\begin{aligned} \underline{R}(a; \mathcal{M}) &= \mathbb{E} \left[ \sum_{t=1}^T (X_t(a, \underline{M}_t^a(a)) - X_t(a, \mathcal{M}(a; \nu_t))) \right] \\ &= \sum_{t=1}^T \mu(a, \underline{M}_t^a(a)) - \mu(a, \mathcal{M}(a; \nu_t)) \\ &= \sum_{t=1}^T \mathbf{1}(\mathcal{M}(\cdot; \nu_t) \in S(\mu)) (\mu(a, \underline{M}_t^a(a)) - \mu(a, \mathcal{M}(a; \nu_t))) \\ &\quad + \sum_{t=1}^T \mathbf{1}(\mathcal{M}(\cdot; \nu_t) \notin S(\mu)) (\mu(a, \underline{M}_t^a(a)) - \mu(a, \mathcal{M}(a; \nu_t))) \\ &\leq \sum_{t=1}^T \mathbf{1}(\mathcal{M}(\cdot; \nu_t) \notin S(\mu)) (\mu(a, \underline{M}_t^a(a)) - \mu(a, \mathcal{M}(a; \nu_t))), \end{aligned}$$

where the last line follows from the definition of  $\underline{M}_t^a$ . Converting the sum over time steps to a sum over possible matches:

$$\begin{aligned} \underline{R}(a; \mathcal{M}) &\leq \sum_{t=1}^T \mathbf{1}(\mathcal{M}(\cdot; \nu_t) \notin S(\mu)) (\mu(a, \underline{M}_t^a(a)) - \mu(a, \mathcal{M}(a; \nu_t))) \\ &= \sum_{M' \notin S(\mu)} T(\mathcal{M}(\cdot; \nu_t) = M') (\mu(a, \underline{M}_t^a(a)) - \mu(a, M'(a; \nu_t))). \end{aligned}$$

Note that, in this setting,  $\underline{M}_t^a(a)$  does not depend on  $t$  because  $\underline{M}_t^a \in \arg \min_{M \in S(\mu)} \mu(a, M(a))$ . Let  $\underline{\Delta}(a, a') = \mu(a, \underline{M}_t^a(a)) - \mu(a, a')$ . Then,

$$\begin{aligned} \underline{R}(a; \mathcal{M}) &\leq \sum_{M' \notin S(\mu)} T(\mathcal{M}(\cdot; \nu_t) = M') \underline{\Delta}(a, M'(a)) \\ &\leq \max_{a' \in \mathcal{A}^+} \underline{\Delta}(a, a') T(\mathcal{M}(\cdot; \nu_t) \notin S(\mu)). \end{aligned}$$

Then, applying Lemma C.6:

$$\begin{aligned} \underline{R}(a; \mathcal{M}) &\leq \max_{a' \in \mathcal{A}} \underline{\Delta}(a, a') T(\mathcal{M}(\cdot; \nu_t) \notin S(\mu)) \\ &\leq 2N^2 L \Delta_{\max}(a) \left( \frac{8\sigma^2 \alpha \log(T)}{\Delta_{\min}^2} + \frac{\alpha}{\alpha - 2} \right). \end{aligned}$$

**Linear optimal regret.** Recall that  $\bar{M}_t^a \in \arg \max_{M \in S(V(\cdot, \cdot; \mu))} \mathbb{E}[U_t(a, M(a); \nu_t)]$ . In the no-cost, no-transfer setting,  $\bar{M}_t^a \in \arg \max_{M \in S(\mu)} \mu(a, M(a))$  such that

$$\begin{aligned} \bar{R}(a; \mathcal{M}) &= \mathbb{E} \left[ \sum_{t=1}^T (X_t(a, \bar{M}_t^a(a)) - X_t(a, \mathcal{M}(a; \nu_t))) \right] \\ &= \sum_{t=1}^T \mu(a, \bar{M}_t^a(a)) - \mu(a, \mathcal{M}(a; \nu_t)) \\ &\leq T \Delta_{\max}(a). \end{aligned} \tag{18}$$

Since  $a' \neq a'' \implies \mu(a, a') \neq \mu(a, a'')$  for all  $a \in \mathcal{A}$ ,  $\Delta_{\max}(a) > 0$ , which implies that  $\bar{R}(a; \mathcal{M}) = O(T)$  for all agents  $a \in \mathcal{A}$ . In other words, the optimal regret of any agent can be upper bounded by a function that grows linearly in  $T$ . To show that this upper bound is tight, we now construct an example in which at least one agent

$a \in \mathcal{A}$  experiences regret that grows  $\Omega(T)$ , i.e., the optimal regret of at least one agent  $a$  can be lower bounded by a function that grows linearly in  $T$  such that  $\bar{R}(a; \mathcal{M}) = \Omega(T)$ .

Suppose  $\mathcal{U} = \{u_1, u_2, u_3\}$ ,  $\mathcal{P} = \{p_1, p_2, p_3\}$ , and

$$\begin{aligned} u_1 \succ_{\mu(p_1, \cdot)} u_2 \succ_{\mu(p_1, \cdot)} u_3, & & p_3 \succ_{\mu(u_1, \cdot)} p_2 \succ_{\mu(u_1, \cdot)} p_1, \\ u_1 \succ_{\mu(p_2, \cdot)} u_2 \succ_{\mu(p_2, \cdot)} u_3, & & p_3 \succ_{\mu(u_2, \cdot)} p_1 \succ_{\mu(u_2, \cdot)} p_2, \\ u_3 \succ_{\mu(p_3, \cdot)} u_2 \succ_{\mu(p_3, \cdot)} u_1, & & p_1 \succ_{\mu(u_3, \cdot)} p_2 \succ_{\mu(u_3, \cdot)} p_3. \end{aligned}$$

Then, the only two stable matchings under  $V(\cdot, \cdot; \mu)$  are

$$M_1(a) = \begin{cases} p_2, & a = u_1, \\ p_1, & a = u_2, \\ p_3, & a = u_3, \\ u_2, & a = p_1, \\ u_1, & a = p_2, \\ u_3, & a = p_3. \end{cases} \quad M_2(a) = \begin{cases} p_2, & a = u_1, \\ p_3, & a = u_2, \\ p_1, & a = u_3, \\ u_3, & a = p_1, \\ u_1, & a = p_2, \\ u_2, & a = p_3. \end{cases}$$

As such,  $\bar{M}_t^{u_1} \in \{M_1, M_2\}$ ,  $\bar{M}_t^{u_2} = M_2$ ,  $\bar{M}_t^{u_3} = M_2$ ,  $\bar{M}_t^{p_1} = M_1$ ,  $\bar{M}_t^{p_2} \in \{M_1, M_2\}$ , and  $\bar{M}_t^{p_3} = M_1$ . The optimal regret of  $p_1$  and  $u_2$  can be lower bounded as follows.

$$\begin{aligned} \bar{R}(a; \mathcal{M}) &= \mathbb{E} \left[ \sum_{t=1}^T (X_t(a, \bar{M}_t^a(a)) - X_t(a, \mathcal{M}(a; \nu_t))) \right] \\ &= \sum_{t=1}^T \mu(a, \bar{M}_t^a(a)) - \mu(a, \mathcal{M}(a; \nu_t)) \\ &\geq T(\mathcal{M}(a; \nu_t) = M_1)(\mu(a, \bar{M}_t^a(a)) - \mu(a, M_1(a))) \\ &\quad + T(\mathcal{M}(a; \nu_t) = M_2)(\mu(a, \bar{M}_t^a(a)) - \mu(a, M_2(a))), \end{aligned} \tag{19}$$

which implies that

$$\begin{aligned} \bar{R}(p_1; \mathcal{M}) &\geq T(\mathcal{M}(\cdot; \nu_t) = M_1)(\mu(p_1, \bar{M}_t^{p_1}(p_1)) - \mu(p_1, M_1(p_1))) \\ &\quad + T(\mathcal{M}(\cdot; \nu_t) = M_2)(\mu(p_1, \bar{M}_t^{p_1}(p_1)) - \mu(p_1, M_2(p_1))) \\ &= T(\mathcal{M}(\cdot; \nu_t) = M_1)(\mu(p_1, M_1(p_1)) - \mu(p_1, M_1(p_1))) \\ &\quad + T(\mathcal{M}(\cdot; \nu_t) = M_2)(\mu(p_1, M_1(p_1)) - \mu(p_1, M_2(p_1))) \\ &= T(\mathcal{M}(\cdot; \nu_t) = M_2)(\mu(p_1, M_1(p_1)) - \mu(p_1, M_2(p_1))) \\ &= T(\mathcal{M}(\cdot; \nu_t) = M_2)(\mu(p_1, u_2) - \mu(p_1, u_3)). \end{aligned} \tag{20}$$

Recall that  $\mu(p_1, u_2) - \mu(p_1, u_3) > 0$ . Therefore, the optimal regret for  $p_1$  grows linearly in  $T$  if  $T(\mathcal{M}(\cdot; \nu_t) = M_2)$  grows linearly in  $T$ . Similarly, applying (19) to agent  $u_2$  gives

$$\begin{aligned} \bar{R}(u_2; \mathcal{M}) &\geq T(\mathcal{M}(\cdot; \nu_t) = M_1)(\mu(u_2, \bar{M}_t^{u_2}(u_2)) - \mu(u_2, M_1(u_2))) \\ &\quad + T(\mathcal{M}(\cdot; \nu_t) = M_2)(\mu(u_2, \bar{M}_t^{u_2}(u_2)) - \mu(u_2, M_2(u_2))) \\ &= T(\mathcal{M}(\cdot; \nu_t) = M_1)(\mu(u_2, M_2(u_2)) - \mu(u_2, M_1(u_2))) \\ &\quad + T(\mathcal{M}(\cdot; \nu_t) = M_2)(\mu(u_2, M_2(u_2)) - \mu(u_2, M_2(u_2))) \\ &= T(\mathcal{M}(\cdot; \nu_t) = M_1)(\mu(u_2, M_2(u_2)) - \mu(u_2, M_1(u_2))) \\ &= T(\mathcal{M}(\cdot; \nu_t) = M_1)(\mu(u_2, p_3) - \mu(u_2, p_1)). \end{aligned} \tag{21}$$

Recall that  $\mu(u_2, p_3) - \mu(u_2, p_1) > 0$ . Therefore, the optimal regret for  $u_2$  grows linearly in  $T$  if  $T(\mathcal{M}(\cdot; \nu_t) = M_1)$  grows linearly in  $T$ .

We now show that either  $T(\mathcal{M}(\cdot; \nu_t) = M_1)$  or  $T(\mathcal{M}(\cdot; \nu_t) = M_2)$  grows linearly in  $T$ , or both do.

$$\begin{aligned}
 T &= \sum_{M \in \mathcal{W}} T(\mathcal{M}(\cdot; \nu_t) = M) \\
 &= \sum_{M \in S(\mu)} T(\mathcal{M}(\cdot; \nu_t) = M) + \sum_{M \notin S(\mu)} T(\mathcal{M}(\cdot; \nu_t) = M), \\
 \implies \sum_{M \in S(\mu)} T(\mathcal{M}(\cdot; \nu_t) = M) &= T - \sum_{M \notin S(\mu)} T(\mathcal{M}(\cdot; \nu_t) = M) \\
 &= T - T(\mathcal{M}(\cdot; \nu_t) \notin S(\mu)) \\
 &\geq T - 2N^2L \left( \frac{8\sigma^2\alpha \log(T)}{\Delta_{\min}^2} + \frac{\alpha}{\alpha - 2} \right)
 \end{aligned}$$

which implies that  $T(\mathcal{M}(\cdot; \nu_t) = M_1) + T(\mathcal{M}(\cdot; \nu_t) = M_2) = \Omega(T)$ . Since  $T(\mathcal{E}) \geq 0$  and  $T > 0$ , it is not possible that  $T(\mathcal{M}(\cdot; \nu_t) = M_1)$  and  $T(\mathcal{M}(\cdot; \nu_t) = M_2)$  both grow sub-linearly in  $T$ . Therefore,  $T(\mathcal{M}(\cdot; \nu_t) = M_1) = \Omega(T)$ ,  $T(\mathcal{M}(\cdot; \nu_t) = M_2) = \Omega(T)$ , or both. By (20) and (21), the optimal regret of at least one agent grows  $\Omega(T)$ .

Intuitively, since it is impossible for the matching at a given time step to simultaneously be both  $M_1$  and  $M_2$ , at least one agent must incur a non-zero optimal regret for  $\Omega(T)$  time steps.

**Fairness not guaranteed.** To show that fairness is not guaranteed, it suffices to show the existence of an example in which the matchings are unfair. That fairness is not guaranteed follows directly from the example under ‘‘Linear optimal regret’’ above. In particular, we use the same example and show that it is possible for at least one agent’s optimal regret to grow  $O(\log(T))$  while another agent’s regret grows  $\Omega(T)$ .

Suppose  $T(\mathcal{M}(\cdot; \nu_t) = M_2) = O(\log(T))$  such that  $T(\mathcal{M}(\cdot; \nu_t) = M_1) = \Omega(T)$ . (This is possible under Algorithm 2, which returns matchings  $\mathcal{M}(\cdot; \nu_t)$  that are  $\mathcal{P}$ -optimal at every time step by Lemma B.3.) As such, for all  $p \in \mathcal{P}$ ,

$$\begin{aligned}
 \bar{R}(p; \mathcal{M}) &= \sum_{t=1}^T \mu(p, \bar{M}_t^p(p)) - \mu(p, \mathcal{M}(p; \nu_t)) \\
 &= (\mu(p, \bar{M}_t^p(p)) - \mu(p, M_1(p)))T(\mathcal{M}(\cdot; \nu_t) = M_1) \\
 &\quad + (\mu(p, \bar{M}_t^p(p)) - \mu(p, M_2(p)))T(\mathcal{M}(\cdot; \nu_t) = M_2) \\
 &\quad + \sum_{M \in \mathcal{W} \setminus \{M_1, M_2\}} (\mu(p, \bar{M}_t^p(p)) - \mu(p, M(p)))T(\mathcal{M}(\cdot; \nu_t) = M) \\
 &= (\mu(p, \bar{M}_t^p(p)) - \mu(p, M_2(p)))T(\mathcal{M}(\cdot; \nu_t) = M_2) \\
 &\quad + \sum_{M \in \mathcal{W} \setminus \{M_1, M_2\}} (\mu(p, \bar{M}_t^p(p)) - \mu(p, M(p)))T(\mathcal{M}(\cdot; \nu_t) = M) \\
 &\leq \Delta_{\max}(a) \left[ T(\mathcal{M}(\cdot; \nu_t) = M_2) + \sum_{M \in \mathcal{W} \setminus \{M_1, M_2\}} T(\mathcal{M}(\cdot; \nu_t) = M) \right] \\
 &= O(\log(T)),
 \end{aligned}$$

where the second to last line follows from the fact that  $\bar{M}_t^p = M_1$  for all  $p \in \mathcal{P}$ , and the last line follows from Lemma C.6 and that  $\mathcal{M}$  can be chosen such that  $T(\mathcal{M}(\cdot; \nu_t) = M_2) = O(\log(T))$ . However, as shown in the example under ‘‘Linear optimal regret’’, when  $T(\mathcal{M}(\cdot; \nu_t) = M_2) = O(\log(T))$ , it must hold true that  $T(\mathcal{M}(\cdot; \nu_t) = M_1) = \Omega(T)$ , which implies that the regret of a user must grow  $\Omega(T)$ . By the definition of fairness (Definition 3), since there exists at least one agent (i.e., a provider) whose optimal regret grows  $O(\log(T))$  and at least one agent (i.e., a user) whose optimal regret grows  $\Omega(T)$ , fairness is not guaranteed.

**High social welfare not guaranteed.** To show that high social welfare is not guaranteed, it suffices to show the existence of an example in which the social welfare is low. Let us again use the example under ‘‘Linear optimal regret’’. Fix  $\kappa > 1$ . Suppose  $\mathcal{M}(\cdot; \nu_t) = M_1$  and  $\nu_t(u_2, p_3) > \frac{1}{\kappa} W_t(\mathcal{M}) = \frac{1}{\kappa} \sum_{a \in \mathcal{A}} \nu_t(a, M_1(a))$ . This outcome is possible because (1)  $u_2$  and  $p_3$  are not matched under  $M_1$ , which means that the values over which

$W_t(\mathcal{M})$  is summed does not include  $\nu_t(u_2, p_3)$ ; (2) the magnitude of  $u_2$ 's transient preferences does not affect the GS algorithm; and (3) it is possible for  $\nu_t(u_2, p_3) > \nu_t(u_2, p_1)$  even when  $\mathcal{M}(u_2; \nu_t) = p_1$  (simply apply the same logic as that given in the example under ‘‘Linear optimal regret’’). As such,  $\max_{M \in \mathcal{W}} \sum_{a \in \mathcal{A}} \nu_t(a, M(a)) \geq \nu_t(u_2, p_3) > \frac{1}{\kappa} W_t(\mathcal{M})$ . Since this reasoning holds true for any choice of  $\kappa > 1$ , by Definition 4, the social welfare is low. Therefore, high social welfare is not guaranteed under no costs and no transfers.  $\square$

**Theorem 2.** *There exist cost and transfer rules  $\mathcal{C}(\cdot, \cdot; \nu_t)$  and  $\mathcal{T}(\cdot, \cdot; \nu_t)$  such that, if the GS algorithm is applied over  $V(\cdot, \cdot; \nu_t)$  at every  $t \in [T]$ , then stability, low regret, fairness, and  $\frac{1}{2}$ -high social welfare are guaranteed for all  $(\mathcal{A}, \mu)$ .*

*Proof.* This is an existence statement. We prove existence in the proof of Theorem 5.  $\square$

## E PROOFS FOR PROPORTIONAL-COST, NO-TRANSFER SETTING (SECTION 5.1)

**Lemma E.7.** *Suppose that  $\mathcal{C}$  and  $\mathcal{T}$  are set according to (1). If  $\gamma = 1$ , then all matchings are stable, i.e., for any  $\psi : \mathcal{A} \times \mathcal{A}^+ \rightarrow \mathbb{R}$  and  $M \in \mathcal{W}$ ,  $M \in S(V(\cdot, \cdot; \psi))$ .*

*Proof.* When  $\gamma = 1$ ,  $V_t(\cdot, \cdot; \psi) = 0$  for all  $\psi : \mathcal{A} \times \mathcal{A}^+ \rightarrow \mathbb{R}$ . Under any matching  $M \in \mathcal{W}$ , there cannot exist another matching  $M' \in \mathcal{W}$  such that  $V(a, M'(a); \psi) > V(a, M(a); \psi)$  for some  $a$  since  $0 > 0$  is always false, which implies that  $M$  is stable. This result holds for any  $M \in \mathcal{W}$ .  $\square$

**Proposition 3.** *Suppose that  $\mathcal{C}$  and  $\mathcal{T}$  are set according to (1). If the GS algorithm is applied over  $V(\cdot, \cdot; \nu_t)$  at every  $t \in [T]$ , then the system is stable. If  $\gamma \in [0, 1)$ ,  $\underline{R}(a; \mathcal{M}) \leq 2N^2L(1 - \gamma)\Delta_{\max}(a) \left( \frac{8\sigma^2\alpha \log(T)}{(1-\gamma)^2\Delta_{\min}^2} + \frac{\alpha}{\alpha-2} \right)$  and, if  $\gamma = 1$ ,  $\underline{R}(a; \mathcal{M}) \leq 0$ . However, under stability, it is not possible to guarantee fairness or high social welfare, and there exist settings  $(\mathcal{A}, \mu)$  for which  $\bar{R}(a; \mathcal{M}) = \Omega(T)$  for at least one agent  $a \in \mathcal{A}$ .*

*Proof.* For  $\gamma \in [0, 1)$ , the setting described in this claim is equivalent to that in Proposition 1, except that the payoffs are scaled by  $(1 - \gamma)$ . To see this, observe that  $V(a, a'; \psi) = \psi(a, a') - \gamma\psi(a, a') = (1 - \gamma)\psi(a, a')$ . Therefore, the same results as in Proposition 1 can be applied here with the appropriate scaling.

We consider  $\gamma = 1$  next. By Lemma B.2, the system is stable since the GS algorithm is applied at each time step. Next, we perform the regret analysis.

**Non-positive pessimal regret.** By (1),  $\mathcal{T} = 0$  and  $\mathcal{C}(a, a'; \nu_t) = \nu_t(a, a')$  for all  $a \in \mathcal{A}$ ,  $a' \in \mathcal{A}^+$ , and  $t \in [T]$  when  $\gamma = 1$ . As such, the observed payoff at each time step  $t$  is  $U_t(a, a'; \nu_t) = X_t(a, a') - \nu_t(a, a')$ , and  $\mathbb{E}[U_t(a, a'; \nu_t)] = \mu(a, a') - \mathbb{E}[\nu_t(a, a')]$ . Recall that

$$\begin{aligned} \underline{M}_t^a &\in \arg \min_{M \in S(V(\cdot, \cdot; \mu))} \mathbb{E}[U_t(a, M(a); \nu_t)] \\ &= \arg \min_{M \in S(V(\cdot, \cdot; \mu))} \mathbb{E}[X_t(a, M(a)) - \nu_t(a, M(a))]. \end{aligned}$$

Under the cost rule,  $V(\cdot, \cdot; \psi) = 0$  for all choices of  $\psi$ . By Lemma E.7,

$$\begin{aligned} \underline{M}_t^a &\in \arg \min_{M \in \mathcal{W}} \mathbb{E}[X_t(a, M(a)) - \nu_t(a, M(a))] \\ &= \arg \min_{M \in \mathcal{W}} (\mu(a, M(a)) - \mathbb{E}[\nu_t(a, M(a))]). \end{aligned}$$

The pessimal regret for  $a$  is

$$\begin{aligned}
 \underline{R}(a; \mathcal{M}) &= \sum_{t=1}^T \mathbb{E} [U_t(a, \underline{M}_t^a(a); \nu_t) - U_t(a, \mathcal{M}(a; \nu_t); \nu_t)] \\
 &= \sum_{t=1}^T \mathbb{E} [X_t(a, \underline{M}_t^a(a)) - \nu_t(a, \underline{M}_t^a(a); \nu_t) - X_t(a, \mathcal{M}(a; \nu_t)) + \nu_t(a; \mathcal{M}(a; \nu_t); \nu_t)] \\
 &= \sum_{t=1}^T \mathbb{E} [X_t(a, \underline{M}_t^a(a)) - \nu_t(a, \underline{M}_t^a(a); \nu_t) - X_t(a, \mathcal{M}(a; \nu_t)) + \nu_t(a; \mathcal{M}(a; \nu_t); \nu_t)].
 \end{aligned}$$

Recall that both domains  $S(V(\cdot, \cdot; \mu))$  and  $S(V(\cdot, \cdot; \nu_t))$  over which  $\underline{M}_t^a$  and  $\mathcal{M}(\cdot; \nu_t)$  are chosen are unrestricted (i.e.,  $S(V(\cdot, \cdot; \mu)) = S(V(\cdot, \cdot; \nu_t)) = \mathcal{W}$ ). Then, by definition of  $\underline{M}_t^a$ ,

$$\mathbb{E} [X_t(a, \underline{M}_t^a(a)) - \nu_t(a, \underline{M}_t^a(a); \nu_t) - X_t(a, \mathcal{M}(a; \nu_t)) + \nu_t(a; \mathcal{M}(a; \nu_t); \nu_t)] \leq 0,$$

which implies  $\underline{R}(a; \mathcal{M}) \leq 0$ , as stated in the result.

**Linear optimal regret.** By definition, the optimal regret of any agent grows  $O(T)$  (see (18) for reference), i.e., the optimal regret of any agent can be upper bounded by a function that grows linearly in  $T$  such that  $\bar{R}(a; \mathcal{M}) = O(T)$  for all  $a \in \mathcal{A}$ . To show that this upper bound is tight, we now construct an example in which at least one agent  $a \in \mathcal{A}$  experiences regret that grows  $\Omega(T)$ , i.e., the optimal regret of at least one agent  $a$  can be lower bounded by a function that grows linearly in  $T$  such that  $\bar{R}(a; \mathcal{M}) = \Omega(T)$ .

$$\begin{aligned}
 \bar{R}(a; \mathcal{M}) &= \sum_{t=1}^T \mathbb{E} [U_t(a, \bar{M}_t^a(a); \nu_t) - U_t(a, \mathcal{M}(a; \nu_t); \nu_t)] \\
 &= \sum_{t=1}^T \mathbb{E} [X_t(a, \bar{M}_t^a(a)) - \nu_t(a, \bar{M}_t^a(a); \nu_t) - X_t(a, \mathcal{M}(a; \nu_t)) + \nu_t(a; \mathcal{M}(a; \nu_t); \nu_t)] \\
 &= \sum_{t=1}^T -\mathbb{E} \left[ \sqrt{\frac{2\sigma^2\alpha \log(t)}{T_{t-1}(a, \bar{M}_t^a(a))}} \right] + \mathbb{E} \left[ \sqrt{\frac{2\sigma^2\alpha \log(t)}{T_{t-1}(a, \mathcal{M}(a; \nu_t))}} \right] \\
 &= \sum_{t=1}^T \sqrt{2\sigma\alpha \log(t)} \mathbb{E} \left[ \sqrt{\frac{1}{T_{t-1}(a, \mathcal{M}(a; \nu_t))}} - \sqrt{\frac{1}{T_{t-1}(a, \bar{M}_t^a(a))}} \right]. \tag{22}
 \end{aligned}$$

Recall that

$$\begin{aligned}
 \bar{M}_t^a &\in \arg \max_{M \in S(V(\cdot, \cdot; \mu))} \mathbb{E} [U_t(a, M(a); \nu_t)] \\
 &= \arg \max_{M \in S(V(\cdot, \cdot; \mu))} \mathbb{E} [X_t(a, M(a)) - \nu_t(a, M(a))] \\
 &= \arg \max_{M \in S(V(\cdot, \cdot; \mu))} -\mathbb{E} \left[ \sqrt{\frac{1}{T_{t-1}(a, M(a))}} \right] \\
 &= \arg \max_{M \in S(V(\cdot, \cdot; \mu))} \mathbb{E} [T_{t-1}(a, M(a))].
 \end{aligned}$$

Under the cost rule,  $V(\cdot, \cdot; \psi) = 0$  for all choices of  $\psi$ , which implies that

$$\bar{M}_t^a \in \arg \max_{M \in \mathcal{W}} \mathbb{E} [T_{t-1}(a, M(a))]. \tag{23}$$

Using this observation, we now construct an example to show that the optimal regret of an agent can grow linearly. To do so, we make repeated use of Lemma E.7, which implies that both of the domains  $S(V(\cdot, \cdot; \mu))$  and  $S(V(\cdot, \cdot; \nu_t))$  over which  $\bar{M}_t^a$  and  $\mathcal{M}(\cdot; \nu_t)$  are chosen are unrestricted (i.e.,  $S(V(\cdot, \cdot; \mu)) = S(V(\cdot, \cdot; \nu_t)) = \mathcal{W}$ ).

Suppose  $p_1, p_2 \in \mathcal{P}$ , where  $N \geq L > 2$ . Suppose that, at every time step,  $\mathcal{M}$  matches  $p_1$  to  $u_1$ ,  $p_2$  uniformly at random to a user  $u \in \mathcal{U} \setminus u_1$ , and all remaining providers arbitrarily to complete the matching. Let  $\eta_t^*(a) =$

$\arg \max_{a' \in \mathcal{A}} \mathbb{E}[T_{t-1}(a, a')]$ . Then, by Lemma E.7, such matchings are stable and can therefore be the matching chosen by the platform at any time step  $t$  to satisfy stability. Under this matching strategy,  $p_1$  is guaranteed to be matched to its optimal match  $\bar{M}_t^{p_1}(p_1)$  at any time step  $t > \max_{u \in \mathcal{U}} T_0(p_1, u)$ , which implies  $\bar{R}(p_1; \mathcal{M}) = O(1)$ . To see this, observe that, after matching  $p_1$  with  $u_1$  for  $\max_{u \in \mathcal{U}} T_0(p_1, u)$  times,  $\eta_t^*(p_1) = u_1$ . In other words, the agent with which  $p_1$  has been most matched is  $u_1$ , which means that it is  $u_1$  is  $p_1$ 's optimal match by (23).

On the other hand, whenever  $|\eta_t^*(p_2)| \neq N$ ,  $p_2$  is not matched to a user  $u \in \eta_t^*(p_2)$  and incurs non-zero regret with probability at least  $1/N$ . The event  $|\eta_t^*(p_2)| = N$  can occur at most  $\lceil T/N \rceil$  times because  $|\eta_t^*(p_2)| = N$  implies that  $p_2$  has sampled all users the same number of times at time step  $t$ . As a result,  $\bar{R}(p_2; \mathcal{M}) = \Omega(T)$ . We have therefore shown that it is not possible to guarantee low optimal regret (i.e., not possible to ensure that  $\bar{R}(a; \mathcal{M}) = O(\log(T))$  for all  $a \in \mathcal{A}$ ) under  $\gamma = 1$ .

Intuitively, because all matchings are stable by Lemma E.7, the platform is afforded so much flexibility that it can choose matchings with detrimental regret outcomes. We now show that this flexibility also allows the platform to choose matchings that are unfair.

**Fairness not guaranteed.** That matchings are not guaranteed to be fair under proportional costs and  $\gamma = 1$  follows directly from the example provided under ‘‘Linear optimal regret’’ above. We showed that, in this example, there exists at least one agent  $a$  for which  $\bar{R}(a; \mathcal{M}) = O(1) = O(\log(T))$  while there exists at least one other agent  $a'$  for which  $\bar{R}(a'; \mathcal{M}) = \omega(\log(T)) = \Omega(T)$ . By Definition 3, this outcome is unfair, showing that fairness is not guaranteed under  $\gamma = 1$ .

**High social welfare not guaranteed.** Under the cost rule,  $V(\cdot, \cdot; \nu_t) = 0$  at all time steps  $t \in [T]$ . Therefore, by Definition 4, the social welfare at every time step is low.  $\square$

**Corollary 4.** *Suppose that  $\mathcal{C}$  and  $\mathcal{T}$  are set according to (1) and that the platform wishes to maximize  $F_t : \mathcal{W} \rightarrow \mathbb{R}$  at every  $t \in [T]$ . Then, when  $\gamma = 1$ , there is no performance cost to stability: i.e.,  $\arg \max_{M \in \mathcal{W}} F_t(M) \in S(V(\cdot, \cdot; \nu_t))$  for all  $t \in [T]$ .*

*Proof.* This result follows from Lemma E.7. Since all matchings are stable at every time step, all matchings are stable, including  $\arg \max_{M \in \mathcal{W}} F_t(M)$ .  $\square$

## F PROOFS FOR BALANCED TRANSFER (SECTION 5.2)

Recall that the GS algorithm in Algorithm 2 is performed with providers as the proposers (i.e., in the notation of Lemma B.3,  $\mathcal{G} = \mathcal{P}$ ). A second version of the GS algorithm in which users are the proposers (i.e.,  $\mathcal{G} = \mathcal{U}$ ) is given Algorithm 3.

In our analysis, we use the following well-known result.

**Lemma F.8.** *Suppose  $V : \mathcal{A} \times \mathcal{A}^+ \rightarrow \mathbb{R}$  denotes a payoff function and  $a' \neq a'' \implies V(a, a') \neq V(a, a'')$  for all  $a \in \mathcal{A}$ . Then, a matching  $M \in \mathcal{W}$  is the unique stable matching under payoffs  $V$  (i.e.,  $M = S(V)$ ) if and only if the GS algorithm performed with providers as proposers and the GS algorithm performed with users as proposers return the same matching  $M$ .*

*Proof.* Denote the matching returned when the GS algorithm is performed with providers as proposers by  $M_1$  and the matching returned when the GS algorithm is performed with users as proposers by  $M_2$ .

To prove the if-and-only-if statement, let's begin with the forward direction: that  $M$  being the unique stable matching implies  $M_1 = M_2$ . This result is trivial because  $M_1$  and  $M_2$  must be stable by Lemma B.1. Therefore, since there is only one stable matching under  $V$ ,  $M_1 = M_2 = M$ .



---

**Algorithm 3:** Gale-Shapley algorithm (with users as proposers)

---

**Input:** Set of agents  $\mathcal{A} = \mathcal{U} \cup \mathcal{P}$ , where the set of users is  $\mathcal{U} = \{u_1, u_2, \dots, u_N\}$ , the set of providers is  $\mathcal{P} = \{p_1, p_2, \dots, p_L\}$ ,  $\mathcal{U} \cap \mathcal{P} = \emptyset$ , and  $N \geq L$ . A payoff function  $V : \mathcal{A} \times \mathcal{A}^+ \rightarrow \mathbb{R}$ .

**Output:** Matching  $M \in S(V) \subset \mathcal{W}$ .

```

1 Initialize matching  $M : \mathcal{A} \rightarrow \mathcal{A}^+$  such that  $M(a) = \emptyset$  for all  $a \in \mathcal{A}$ ;
2 Initialize empty (FIFO) queues  $Q(u) = []$  for all  $u \in \mathcal{U}$ ;

// Fill each user's queue with providers in order of decreasing preference.
3 for  $u \in \mathcal{U}$  do
4   for  $i = 1, 2, \dots, L$  do
5     Append  $r^{-1}(i; V(u, \cdot))$  to  $Q(u)$ ; // Add  $u$ 's  $i$ -th ranked user.
6   end
7 end

// As long as there exists a user who is unmatched and has a non-empty queue...
8 while  $\exists u \in \mathcal{U} : M(u) = \emptyset \cap |Q(u)| > 0$  do
9    $p \leftarrow \text{pop}(Q(u))$ ; // User  $u$ 's favorite provider of those remaining in  $u$ 's queue.

// If provider  $p$  is unmatched, match  $u$  and  $p$ .
10  if  $M(p) = \emptyset$  then
11     $M(p) \leftarrow u$ ;
12     $M(u) \leftarrow p$ ;

// If provider  $p$  prefers  $u$  to its current match  $M(p)$ , match  $u$  and  $p$ .
13  else if  $V(p, u) > V(p, M(p))$  then
14     $M'(M(p)) = \emptyset$ ;
15     $M(p) \leftarrow u$ ;
16     $M(u) \leftarrow p$ ;
17 end
18 Return  $M$ ;
```

---

In the GS algorithm presented here, the users are the proposers. The version of the algorithm in which providers are the proposers is given in Section 4.

To show the backward direction, we must show that  $M_1 = M_2 = M$  implies that  $M$  is the unique stable matching under  $V$ . By Lemma B.3,  $M_1$  is  $\mathcal{P}$ -optimal and  $\mathcal{U}$ -pessimal, while  $M_2$  is  $\mathcal{U}$ -optimal and  $\mathcal{P}$ -pessimal. Since  $M_1 = M_2 = M$ , applying Definition B.1 gives that:

$$\begin{aligned}
V(p, M(p)) &\geq V(p, M'(p)), \\
V(p, M(p)) &\leq V(p, M'(p)), \\
V(u, M(u)) &\leq V(u, M'(u)), \\
V(u, M(u)) &\geq V(u, M'(u)),
\end{aligned}$$

for all  $p \in \mathcal{P}$ ,  $u \in \mathcal{U}$ , and  $M' \in S(V)$ . However, the first two statements imply that  $V(p, M(p)) \geq V(p, M'(p)) \geq V(p, M(p))$  for all  $p \in \mathcal{P}$  and all stable matchings  $M' \in S(V)$ . Similarly, the latter two statements imply that  $V(u, M(u)) \leq V(u, M'(u)) \leq V(u, M(u))$  for all  $u \in \mathcal{U}$  and all stable matchings  $M' \in S(V)$ . These outcomes are only possible if  $M' = M$  for all  $M' \in S(V)$ , which implies that there is a unique stable matching  $M = M_1 = M_2$ .  $\square$

**Definition 5.** A function  $h : \mathcal{A} \times \mathcal{A}^+ \rightarrow \mathbb{R}$  is pairwise-unique if  $h(a_1, a_2) \neq h(a_3, a_4)$  unless  $(a_1, a_2) = (a_3, a_4)$  or  $(a_2, a_1) = (a_3, a_4)$ .

Recall that we assume pairwise-uniqueness in Section 3.1. In the Appendix, in the remainder of this section, if pairwise-uniqueness is needed, it is mentioned explicitly. Outside of this section, it is not needed, except for Theorem 2. However, as noted in the main text, pairwise-uniqueness is not strictly necessary for any of the results. Instead, the platform could have a consistent tie-breaking rule. Specifically, the platform would first place matches (i.e., user-provider pairs) in some order. This order could be principled or arbitrary. Then, whenever there is a tie, the algorithm breaks ties based on this order. The same stability, regret, fairness, and social welfare results would hold with small modifications. In particular, the set of stable matchings in this scenario should be restricted to stable matchings that can be obtained using this tie-breaking rule.

---

**Algorithm 4:** Greedy algorithm to determine a stable matching under the balanced transfer rule.

---

**Input:** Set of agents  $\mathcal{A} = \mathcal{U} \cup \mathcal{P}$ , where the set of users is  $\mathcal{U} = \{u_1, u_2, \dots, u_N\}$ , the set of providers is  $\mathcal{P} = \{p_1, p_2, \dots, p_L\}$ ,  $\mathcal{U} \cap \mathcal{P} = \emptyset$ , and  $N \geq L$ . Preferences  $\psi : \mathcal{A} \times \mathcal{A}^+ \rightarrow \mathbb{R}$ .

**Output:** Matching  $M \in S(V) \subset \mathcal{W}$ , where  $V(a, a') = \frac{1}{2}(\psi(a, a') + \psi(a', a))$ .

```

1 Initialize matching  $M : \mathcal{A} \rightarrow \mathcal{A}^+$  such that  $M(a) = \emptyset$  for all  $a \in \mathcal{A}$ ;
2  $w(e) \leftarrow \psi(u, p) + \psi(p, u)$  for all  $e = (u, p)$ , where  $u \in \mathcal{U}$  and  $p \in \mathcal{P}$ ;
3 Fill a (FIFO) queue  $\bar{Q}$  with edges  $e$  in decreasing order of weight  $w(e)$ ;
   // As long as there exists a provider who is unmatched...
4 while  $\exists p' \in \mathcal{P} : M(p') = \emptyset$  do
5    $e = (u, p) \leftarrow \text{pop}(\bar{Q})$ ; // Edge with highest weight that remains in the queue.
   // If edge  $e$  does not conflict with an existing match, then add it to  $M$ 
6   if  $M(u) = \emptyset$  and  $M(p) = \emptyset$  then
7      $M(u) \leftarrow p$ ;
8      $M(p) \leftarrow u$ ;
9   end
10 end
11 Return  $M$ ;
```

---

**Lemma F.9.** *Suppose that  $\mathcal{C}$  and  $\mathcal{T}$  are set according to (2)-(3). If payoffs  $V(\cdot, \cdot; \psi)$  are pairwise-unique, then there is a unique stable matching under these payoffs: i.e.,  $|S(V(\cdot, \cdot; \psi))| = 1$ .*

*Proof.* Under (2),  $V(a, a'; \psi) = \frac{1}{2}(\psi(a, a') + \psi(a', a)) = V(a', a; \psi)$ . In other words,  $V(\cdot, \cdot; \psi)$  is symmetric. Graphically, this means that the edges between users and providers are undirected. If  $V(\cdot, \cdot; \psi)$  is pairwise-unique, then all edge weights in this undirected graph are unique. The unique stable matching can be found using a greedy algorithm, as given in Algorithm 4. In this algorithm, all edges are first sorted from highest to lowest weight. Then, beginning with the edges with the highest weights, the match  $(u, p)$  associated with the given edge is added to the matching  $M$  unless it conflicts with an existing match in  $M$ . Repeat this process until all providers are matched such that  $M(p) \neq \emptyset$  for all  $p \in \mathcal{P}$ .

To see why the resulting matching  $M$  is stable, observe that it is identical to the GS algorithm in Algorithm 2 as long as the unmatched provider  $p$  who proposes in Line 8 of Algorithm 2 is always the unmatched provider with the greatest unproposed edge weight such that  $\max_{u \in Q(p)} V(p, u; \psi) \geq \max_{u \in Q(p')} V(p', u; \psi)$  for all  $p' \in \mathcal{P}$  for which  $M(p') = \emptyset$ . This is equivalent to the greedy algorithm described above because it is not possible for any proposed match to override a previous one. As such, all proposals that conflict with an existing match are discarded. As the GS algorithm is agnostic to the order of proposers, such an ordering is allowed, which implies that  $M$  is stable by Lemma B.1.

To see why the resulting stable matching is unique, we make use of Lemma F.8. In particular, we have already shown that  $M$  is the result of running the GS algorithm with providers as proposers. If we additionally show that  $M$  can be obtained by running the GS algorithm with users as proposers, then  $M$  is the unique stable matching under  $V(\cdot, \cdot; \psi)$ . We use the same observation as above.  $M$  can be obtained by running Algorithm 3, in which the user  $u$  chosen to propose in Line 8 is the user with the greatest unproposed edge weight such that  $\max_{p \in Q(u)} V(u, p; \psi) \geq \max_{p \in Q(u')} V(u', p; \psi)$  for all  $u' \in \mathcal{U}$  for which  $M(u') = \emptyset$  and  $|Q(u')| > 0$ . However, this is precisely the same algorithm as that described in the paragraph directly above because  $V(a, a'; \psi) = V(a', a; \psi)$  is symmetric. Therefore, Algorithms 2 and 3 both return  $M$ , which implies that  $M$  is the unique stable matching by Lemma F.8.  $\square$

Let  $\rho(a, a') = \frac{1}{2}(\mu(a, a') + \mu(a', a))$  and  $\Delta_{\min}^\rho = \min_{a_1 \in \mathcal{A}, a_2 \in \mathcal{A}^+, a_3 \in \mathcal{A}^+ \setminus a_2} |\rho(a_1, a_2) - \rho(a_1, a_3)|$ . Note that, if  $\rho$  is pairwise-unique,  $S(\rho)$  contains only one element, which we denote by  $M^{*,\rho} = S(\rho)$ . Accordingly, let  $\Delta_{\max}^{*,\rho}(a) = \max_{a' \in \mathcal{A}^+} (\rho(a, M^{*,\rho}(a)) - \rho(a, a'))$ .

**Theorem 5.** *Suppose that  $\mathcal{C}$  and  $\mathcal{T}$  are set according to (2)-(3). Then,  $|S(\rho)| = 1$ . Let  $M^{*,\rho}$  be the only element in  $S(\rho)$  and  $\Delta_{\max}^{*,\rho}(a) = \max_{a' \in \mathcal{A}^+} (\rho(a, M^{*,\rho}(a)) - \rho(a, a'))$ . If the GS algorithm is applied over  $V(\cdot, \cdot; \nu_t)$  at every  $t \in [T]$ , then the system is stable, the social welfare is  $\frac{1}{2}$ -high at all  $t \in [T]$ ,  $\mathcal{M}$  is fair, and  $\underline{R}(a; \mathcal{M}) = \bar{R}(a; \mathcal{M}) \leq \Delta_{\max}^{*,\rho}(a) N^2 L \left( \frac{8\sigma^2 \alpha \log(T)}{(\Delta_{\min}^\rho)^2} + \frac{\alpha}{\alpha-2} \right)$  for all  $a \in \mathcal{A}$ .*

*Proof.* By Lemma B.2, the system is stable since the GS algorithm is applied at each time step. Next, we perform the regret analysis.

**Logarithmic optimal regret.** By (2),  $\mathcal{C} = 0$  and  $\mathcal{T}(a, a'; \nu_t) = \frac{1}{2}(\nu_t(a', a) - \nu_t(a, a'))$  for all  $a, a' \in \mathcal{A}$  and  $t \in [T]$ . As such, the observed payoff at each time step  $t$  is  $U_t(a, a'; \nu_t) = X_t(a, a') + \frac{1}{2}(\nu_t(a', a) - \nu_t(a, a'))$ , and

$$\begin{aligned} \mathbb{E}[U_t(a, a'; \nu_t)] &= \mu(a, a') + \frac{1}{2}\mathbb{E}[\nu_t(a', a) - \nu_t(a, a')] \\ &= \frac{1}{2}(\mu(a, a') + \mu(a', a)) + \mathbb{E}\left[\sqrt{\frac{2\sigma^2\alpha \log(T)}{T_{t-1}(a', a)}} - \sqrt{\frac{2\sigma^2\alpha \log(T)}{T_{t-1}(a, a')}}\right] \\ &= \frac{1}{2}(\mu(a, a') + \mu(a', a)), \end{aligned} \tag{24}$$

where the last line follows from the fact that the number of times agent  $a$  samples  $a'$  is the same as the number of times agent  $a'$  samples  $a$  since samples are obtained from being matched, which means that  $T_{t-1}(\cdot, \cdot)$  is therefore symmetric.

Recall that  $\underline{R}(a; \mathcal{M}) \leq \bar{R}(a; \mathcal{M})$ . We therefore focus on upper bounding the latter. Under the given setting,  $V(\cdot, \cdot; \psi) = \psi(a, a') + \frac{1}{2}(\psi(a', a) - \psi(a, a')) = \frac{1}{2}(\psi(a', a) + \psi(a, a'))$  for all choices of  $\psi$ . As such,

$$\begin{aligned} \bar{M}_t^a &\in \arg \max_{M \in \mathcal{S}(V(\cdot, \cdot; \mu))} \mathbb{E}[U_t(a, M(a); \nu_t)] \\ &= \arg \max_{M \in \mathcal{S}(V(\cdot, \cdot; \mu))} \frac{1}{2}(\mu(a, M(a)) + \mu(M(a), a)) \\ &= \arg \max_{M \in \mathcal{S}(\rho)} \frac{1}{2}\rho(a, M(a)), \end{aligned} \tag{25}$$

where  $\rho(a, a') = \frac{1}{2}(\mu(a, a') + \mu(a', a))$ , and the second equality follows from (24). By assumption (see Section 3.1),  $\rho$  is pairwise-unique. Applying Lemma F.9,  $|\mathcal{S}(\rho)| = 1$ . In other words, there is a unique stable matching under  $\rho$ . Let  $M^{*\cdot\rho}$  denote this stable matching. Since both optimal and pessimal regret are from the same set  $\mathcal{S}(\rho)$  and  $|\mathcal{S}(\rho)| = 1$ ,  $\bar{M}_t^a = \underline{M}_t^a = M^{*\cdot\rho}$ , and  $\underline{R}(a; \mathcal{M}) = \bar{R}(a; \mathcal{M})$  for all  $a \in \mathcal{A}$  and all  $t \in [T]$ .

By (24),

$$\begin{aligned} \bar{R}(a; \mathcal{M}) &= \sum_{t=1}^T \mathbb{E}[U_t(a, \bar{M}_t^a(a); \nu_t) - U_t(a, \mathcal{M}(a; \nu_t); \nu_t)] \\ &= \sum_{t=1}^T \frac{1}{2}\mathbb{E}[\mu(a, \bar{M}_t^a(a)) + \mu(\bar{M}_t^a(a), a) - \mu(a, \mathcal{M}(a; \nu_t)) - \mu(\mathcal{M}(a; \nu_t), a)] \\ &= \sum_{t=1}^T \frac{1}{2}(\mu(a, M^{*\cdot\rho}(a)) + \mu(M^{*\cdot\rho}(a), a) - \mathbb{E}[\mu(a, \mathcal{M}(a; \nu_t)) + \mu(\mathcal{M}(a; \nu_t), a)]) \\ &= \sum_{t=1}^T \frac{1}{2}(\rho(a, M^{*\cdot\rho}(a)) - \mathbb{E}[\rho(a, \mathcal{M}(a; \nu_t))]) \\ &= \frac{1}{2} \sum_{M \in \mathcal{W}} \mathbb{E}[T(\mathcal{M}(\cdot; \nu_t) = M(\cdot))] (\rho(a, M^{*\cdot\rho}(a)) - \rho(a, M(a))) \\ &\leq \frac{1}{2} \sum_{M \in \mathcal{W} \setminus M^{*\cdot\rho}} \mathbb{E}[T(\mathcal{M}(\cdot; \nu_t) = M(\cdot))] \Delta^*(a, M(a)) \\ &\leq \frac{1}{2} \max_{a' \in \mathcal{A}^+} \Delta^*(a, a') \mathbb{E}[T(\mathcal{M}(\cdot; \nu_t) \neq M^{*\cdot\rho}(\cdot))], \end{aligned} \tag{26}$$

where  $\Delta^*(a, a') = \rho(a, M^{*\cdot\rho}(a)) - \rho(a, a')$ . Similarly to the proof in the no-cost, no-transfer setting, we would like to apply Lemma C.6, but we first need to transform the current setting into an equivalent setting that meets the conditions of Lemma C.6.

To do so, consider a new setting  $(\mathcal{A}, \mu')$ , in which there is no cost or transfer. In this new setting, let the true and transient preferences at time  $t$  be denoted by  $\mu'$ , where  $\mu'(a, a') = \frac{1}{2}(\mu(a, a') + \mu(a', a))$ . All other quantities are defined analogously with respect to  $\mu'$  (e.g.,  $\nu'_t$  is defined as in Section 3.1 with respect to  $\mu'$ ,  $X'_t(a, a')$  are drawn i.i.d. from a  $\sigma^2$ -sub-Gaussian distribution centered at  $\mu'(a, a')$ ). Furthermore, let  $\Delta_{\min}^\rho = \min_{a_1, a_2, a_3 \in \mathcal{A}} |\mu'(a_1, a_2) - \mu'(a_1, a_3)|$ . This is indeed an equivalent setting to that stated in Lemma C.6, except that  $U'_t \neq X'_t$ . However, as evident in the proof of Lemma C.6,  $U'_t$  is not relevant to the proof. Moreover, recall that agents observe their reward  $X_t$  at each time step and know how much is transferred. Agents can therefore distinguish between their reward and transfer. Combining these two facts, Lemma C.6 can be applied because all relevant quantities meet the lemma's conditions. In particular, the only two quantities of interest are  $V(a, a'; \mu) = \frac{1}{2}(\mu(a, a') + \mu(a', a)) = V'(a, a'; \mu')$  and  $V(a, a'; \nu_t) = \frac{1}{2}(\nu_t(a, a') + \nu_t(a', a)) = V'(a, a'; \nu'_t)$ , as required.

Then, applying Lemma C.6 to the equivalent alternate setting  $(\mathcal{A}, \mu')$ :

$$\mathbb{E}T'(M(\cdot; \nu'_t) \notin S(\mu')) = \mathbb{E}[T(\mathcal{M}(\cdot; \nu_t) \neq M^{*,\rho}(\cdot))] \leq 2N^2L \left( \frac{8\sigma^2\alpha \log(T)}{(\Delta_{\min}^\rho)^2} + \frac{\alpha}{\alpha - 2} \right). \quad (27)$$

By definition of  $\mu'$ ,  $\Delta_{\min}^\rho = \frac{1}{2} \min_{a_1 \in \mathcal{A}, a_2 \in \mathcal{A}^+, a_3 \in \mathcal{A}^+ \setminus a_2} |\mu(a_1, a_2) + \mu(a_2, a_1) - \mu(a_1, a_3) - \mu(a_3, a_1)|$ . Combining (27) with (26) yields the result on optimal regret.

**Fairness is guaranteed.** That  $\mathcal{M}$  as described in the result also ensures fairness (i.e., stability guarantees fairness) follows directly from the analysis under ‘‘Logarithmic optimal regret’’ above. Specifically, since  $\bar{R}(a; \mathcal{M}) = O(\log(T))$  for all  $a \in \mathcal{A}$ , fairness is achieved for all agents by Definition 3.

**High social welfare is guaranteed.** That  $\mathcal{M}$  also results in high social welfare (i.e., stability guarantees high social welfare) follows directly from the fact that the greedy algorithm in Algorithm 4 returns a matching whose summed edge weights is within factor-2 of the maximal bipartite weighted matching.

Recall that the social welfare of  $(\mathcal{M}, \mathcal{C}, \mathcal{T})$  at time  $t$  is defined as  $W_t(\mathcal{M}) = \sum_{a \in \mathcal{A}} V(a, \mathcal{M}(a; \nu_t); \nu_t)$ . Under (2),  $V(a, a'; \nu_t) = \frac{1}{2}(\nu_t(a, a') + \nu_t(a', a))$ . Recall from above that applying Algorithm 4 to preferences  $\nu_t$  returns a stable matching under payoffs  $V(\cdot, \cdot; \nu_t)$ . It remains to show that the social welfare obtained by applying Algorithm 4 to preferences  $\nu_t$  is within factor-2 of maximum social welfare possible under any feasible (not necessarily stable) matching in  $\mathcal{W}$  for all  $t \in [T]$ .

Consider Algorithm 4. In the remainder of the proof,  $e = (u, p)$ ,  $e' = (u', p')$ , and  $e'' = (u'', p'')$ . If an edge  $e = (u, p)$  is added by the algorithm, then it does not conflict with any of the edges added before it. In other words, there are no edges  $e' = (u', p')$  for which  $w(e') > w(e)$  such that  $u' = u$  or  $p' = p$ . Applying this observation, we find that, for any  $M \in \mathcal{W}$ ,

$$\begin{aligned} \sum_{a \in \mathcal{A}} V(a, M(a); \nu_t) &= \sum_{a \in \mathcal{A}} \frac{1}{2} (\nu_t(a, M(a)) + \nu_t(M(a), a)) \\ &= \sum_{e: M(u)=p} (\nu_t(u, p) + \nu_t(p, u)) \\ &= \sum_{e: M(u)=p} w(e) \\ &\leq \sum_{e: M(u)=p} \max_{e' \neq e'', u'=u \cap p''=p} (w(e') + w(e'')) \\ &\leq \sum_{e: M(u)=p} \left( \max_{e': u'=u} w(e') + \max_{e'': p''=p} w(e'') \right) \\ &\leq \sum_{e: M(u)=p} (w((u, \mathcal{M}(u; \nu_t))) + w((\mathcal{M}(p; \nu_t), p))) \\ &\leq \sum_{u \in \mathcal{U}} w((u, \mathcal{M}(u; \nu_t))) + \sum_{p \in \mathcal{P}} w((\mathcal{M}(p; \nu_t), p)) \end{aligned}$$

$$\begin{aligned}
 &= \sum_{a \in \mathcal{A}} \nu_t(a, \mathcal{M}(a; \nu_t)) + \nu_t(\mathcal{M}(a; \nu_t), a) \\
 &= 2 \sum_{a \in \mathcal{A}} V(a, \mathcal{M}(a; \nu_t); \nu_t) \\
 &= 2W_t(\mathcal{M}).
 \end{aligned}$$

Since  $\sum_{a \in \mathcal{A}} V(a, M(a); \nu_t) \leq 2W_t(\mathcal{M})$  for all  $M \in \mathcal{W}$ ,  $\max_{M \in \mathcal{W}} \sum_{a \in \mathcal{A}} V(a, M(a); \nu_t) \leq 2W_t(\mathcal{M})$  for all  $t \in [T]$ , which concludes the proof.  $\square$

## G PROOFS FOR PRICING SETTING (SECTION 5.3)

**Lemma G.10.** *Suppose that  $\mathcal{C}$  and  $\mathcal{T}$  are set according to (4)-(5). Suppose there exists a  $B \in \mathbb{R}_{>0}$  such that  $|\mu(u, \cdot)| \leq B$  for all  $u \in \mathcal{U}$ . For any arbitrary ordering of providers  $(p_1, p_2, \dots, p_L)$ , let  $c_1 = 0$ ,  $c_2 = 2B(1 - L)$ , and  $g(p_k; \psi) = 2B(L - k)$  for all  $p_k \in \mathcal{P}$  and  $\psi : \mathcal{A} \times \mathcal{A}^+ \rightarrow \mathbb{R}$ . Then,  $|S(V(\cdot, \cdot; \mu))| = 1$ , i.e., there is a unique stable matching under preferences  $\mu$ .*

*Proof.* By Lemma F.8, there is a unique stable matching under payoffs  $V : \mathcal{A} \times \mathcal{A}^+ \rightarrow \mathbb{R}_{\geq 0}$  where  $a' \neq a'' \implies V(a, a') \neq V(a, a'') \forall a \in \mathcal{A}$  if performing the GS algorithm with providers as proposers and with users as proposers return the same matching  $M^{*,B}$ . We now show that this is indeed the case for  $c_1 = 0$ ,  $c_2 = 2B(1 - L)$ , and  $g(p_k; \psi) = 2B(L - k)$ .

Recall that the GS algorithm with providers and users as proposers can be found in Algorithms 2 and 3, respectively. Suppose that  $M^{*,B}$  is returned by Algorithm 2. Recall that  $M^{*,B}$  is agnostic to which proposer is chosen in Line 8 as long as the proposer is unmatched. Similarly, Algorithm 3 is agnostic to which user is chosen in Line 8 as long as the user is unmatched and as a non-empty queue. To show that Algorithm 3 returns  $M^{*,B}$ , we show that there is a way to pick the user that is proposing in Line 8 such that the resulting matching is  $M^{*,B}$ .

First, note that for any user  $u \in \mathcal{U}$ , we have that:  $V(u, p_k; \mu) = \mu(u, p_k) - g(p_k; \mu) = \mu(u, p_k) - 2B(L - k)$ , which implies that  $V(u, p_k; \mu) - V(u, p_{k-1}; \mu) = \mu(u, p_k) - \mu(u, p_{k-1}) - 2B(L - k) + 2B(L - k + 1) \geq -B + 2B = B > 0$ . By induction,  $V(u, p_k; \mu) > V(u, p_j; \mu)$  for all  $k > j$  and all  $u \in \mathcal{U}$ . That is, all users have the same ordinal preferences over providers:  $p_L \succ_{V(u, \cdot; \mu)} p_{L-1} \succ_{V(u, \cdot; \mu)} \dots \succ_{V(u, \cdot; \mu)} p_1$  for all  $u \in \mathcal{U}$ .

We begin with Algorithm 3. Let user  $M^{*,B}(p_L)$  be the first to propose.  $M^{*,B}(p_L)$  proposes to  $p_L$  since  $p_L$  is the top preference of all users.  $p_L$  accepts since it is unmatched. Then, let  $M^{*,B}(p_{L-1})$  propose twice. It first proposes to  $p_L$  and is rejected (if not, then  $p_L$  prefers  $M^{*,B}(p_{L-1})$  over  $M^{*,B}(p_L)$ , and we already know that all users prefer  $p_L$  the most, which means that  $M^{*,B}$  is not stable, which is a contradiction). It then proposes to  $p_{L-1}$  who accepts since it is unmatched. If we proceed in such a manner, we eventually get to a point in the while loop of Algorithm 3 at which all users and providers are matched according to  $M^{*,B}$ . However, the while loop has not terminated. Even so, the matching will not change because, if a user  $u$  proposes after this point in the while loop, the  $p$  to which  $u$  proposes must be matched.  $u$  therefore prefers to be matched to  $p$  than to be unmatched, but  $p$  cannot prefer to be matched to  $u$  than to  $M^{*,B}(p)$  because such a result would imply that  $M^{*,B}$  is not stable, which is a contradiction. Therefore, the while loop continues until termination, and Algorithm 3 returns  $M^{*,B}$ . Since both Algorithm 2 and 3 return the same matching  $M^{*,B}$ ,  $M^{*,B}$  is the unique stable matching by Lemma F.8.  $\square$

**Lemma G.11.** *Suppose that  $\mathcal{C}$  and  $\mathcal{T}$  are set according to (4)-(5). Suppose there exists a  $B \in \mathbb{R}_{>0}$  such that  $|\mu(u, \cdot)| \leq B$  for all  $u \in \mathcal{U}$ . Then, there always exist constants  $c_1, c_2$  and pricing  $g$  such that  $|S(V(\cdot, \cdot; \mu))| = 1$ , i.e., there is a unique stable matching under payoffs  $V(\cdot, \cdot; \mu)$ .*

*Proof.* This result follows directly from Lemma G.10, which gives one set of costs  $c_1, c_2$  and pricing rule  $g$  under which  $|S(V(\cdot, \cdot; \mu))| = 1$ .  $\square$

**Theorem 6.** *Suppose that  $\mathcal{C}$  and  $\mathcal{T}$  are set according to (4)-(5). If the GS algorithm is applied over  $V(\cdot, \cdot; \nu_t)$  at every  $t \in [T]$ , then the system is stable. Moreover, if there exists a  $B \in \mathbb{R}_{>0}$  such that  $|\mu(u, \cdot)| \leq B$  for all  $u \in \mathcal{U}$ , then  $|S(V(\cdot, \cdot; \mu))| = 1$  and there exist constants  $c_1$  and  $c_2$  as well as a pricing rule  $g$  such that  $\mathcal{M}$  is fair and  $\underline{R}(a; \mathcal{M}) = \bar{R}(a; \mathcal{M}) \leq 2\Delta_{\max}^{*,B}(a)N^2L \left( \frac{8\sigma^2\alpha \log(T)}{(\Delta_{\min})^2} + \frac{\alpha}{\alpha-2} \right)$  for all  $a \in \mathcal{A}$ , where  $M^{*,B}$  is the only element in  $S(V(\cdot, \cdot; \mu))$  and  $\Delta_{\max}^{*,B}(a) = 2B(L - 1)\mathbf{1}(a \in \mathcal{U}) + \max_{a' \in \mathcal{A}^+} (\mu(a, M^{*,B}(a)) - \mu(a, a'))$ .*

*Proof.* By Lemma B.2, the system is stable since the GS algorithm is applied at each time step. Next, we perform the regret analysis.

**Logarithmic optimal regret.** By (4)-(5),  $\mathcal{C}(p, u; \nu_t) = c_1$ ,  $\mathcal{C}(u, p; \nu_t) = c_2$ ,  $\mathcal{T}(p, u; \nu_t) = g(p; \nu_t)$ , and  $\mathcal{T}(u, p; \nu_t) = -g(p; \nu_t)$  for all  $u \in \mathcal{U}$ ,  $p \in \mathcal{P}$ , and  $t \in [T]$ . As such, the observed payoffs at each time step  $t$  are  $U_t(p, u; \nu_t) = X_t(p, u) - c_1 + g(p; \nu_t)$  and  $U_t(u, p; \nu_t) = X_t(u, p) - c_2 - g(p; \nu_t)$  for all  $u \in \mathcal{U}$  and  $p \in \mathcal{P}$ . In addition,  $\mathbb{E}[U_t(p, u; \nu_t)] = \mu(p, u) - c_1 + \mathbb{E}[g(p; \nu_t)]$ ,  $\mathbb{E}[U_t(u, p; \nu_t)] = \mu(u, p) - c_2 - \mathbb{E}[g(p; \nu_t)]$ ,  $V(p, u; \psi) = \psi(p, u) - c_1 + g(p; \psi)$ , and  $V(u, p; \psi) = \psi(u, p) - c_2 - g(p; \psi)$ .

Recall that our goal is to show that there exists a pricing rule  $g$  such that both pessimal and optimal regret are upper bounded by  $O(\log(T))$ . To show existence, let  $c_1 = 0$ ,  $c_2 = 2B(1 - L)$ ,  $g(p_k; \psi) = 2B(L - k)$  for all  $k \in [L]$  and preferences  $\psi$ . By Lemma G.10,  $|S(V(\cdot, \cdot; \mu))| = 1$ , which implies that  $\bar{M}_t^a = \underline{M}_t^a$ . We denote this unique stable matching under payoffs  $V(\cdot, \cdot; \mu)$  by  $M^{*,B}$ . Since  $|S(V(\cdot, \cdot; \mu))| = 1$ ,  $\underline{R}(a; \mathcal{M}) = \bar{R}(a; \mathcal{M})$ . We focus on upper bounding the latter. For any provider  $p_k \in \mathcal{P}$ ,

$$\begin{aligned}
 \bar{R}(p; \mathcal{M}) &= \sum_{t=1}^T (\mu(p, M^{*,B}(p)) - c_1 + \mathbb{E}[g(M^{*,B}(p); \nu_t)] - \mu(p, \mathcal{M}(p; \nu_t)) + c_1 - \mathbb{E}[g(\mathcal{M}(p; \nu_t); \nu_t)]) \\
 &= \sum_{t=1}^T (\mu(p, M^{*,B}(p)) - \mu(p, \mathcal{M}(p; \nu_t)) + 2B(L - k) - 2B(L - k)) \\
 &= \sum_{M \in \mathcal{W}} \mathbb{E}T(\mathcal{M}(\cdot; \nu_t) = M(\cdot)) (\mu(p, M^{*,B}(p)) - \mu(p, M(p))) \\
 &\leq \Delta_{\max}^{*,B}(p) \sum_{M \neq M^{*,B}} \mathbb{E}T(\mathcal{M}(\cdot; \nu_t) = M(\cdot)) \\
 &\leq \Delta_{\max}^{*,B}(p) \mathbb{E}T(\mathcal{M}(\cdot; \nu_t) \neq M^{*,B}(\cdot)), \tag{28}
 \end{aligned}$$

where  $\Delta_{\max}^{*,B}(p) = \max_{u \in \mathcal{U}} (\mu(p, M^{*,B}(p)) - \mu(p, u))$ . For all users  $u \in \mathcal{U}$ ,

$$\begin{aligned}
 \bar{R}(u; \mathcal{M}) &= \sum_{t=1}^T (\mu(u, M^{*,B}(u)) - c_2 - \mathbb{E}[g(M^{*,B}(u); \nu_t)] - \mu(u, \mathcal{M}(u; \nu_t)) + c_2 + \mathbb{E}[g(\mathcal{M}(u; \nu_t); \nu_t)]) \\
 &= \sum_{t=1}^T (\mu(u, M^{*,B}(u)) - \mathbb{E}[g(M^{*,B}(u); \nu_t)] - \mu(u, \mathcal{M}(u; \nu_t)) + \mathbb{E}[g(\mathcal{M}(u; \nu_t); \nu_t)]) \\
 &\leq \Delta_{\max}^{*,B}(u) \sum_{M \neq M^{*,B}} \mathbb{E}T(\mathcal{M}(\cdot; \nu_t) = M(\cdot)) \\
 &= \Delta_{\max}^{*,B}(u) \mathbb{E}T(\mathcal{M}(\cdot; \nu_t) \neq M^{*,B}(\cdot)), \tag{29}
 \end{aligned}$$

where  $\Delta_{\max}^{*,B}(u) = 2B(L - 1) + \max_{p \in \mathcal{P}} (\mu(u, M^{*,B}(u)) - \mu(u, p))$ .

We therefore would like to upper bound  $\mathbb{E}T(\mathcal{M}(\cdot; \nu_t) \neq M^{*,B}(\cdot))$ . Recall that  $\mathcal{M}(\cdot; \nu_t) = S(V(\cdot, \cdot; \nu_t))$  and

$M^{*,B} \in S(V(\cdot, \cdot; \mu))$ . Then, following approximately the same steps as in the proof in Lemma C.6:

$$\begin{aligned}
 \mathbb{E}T(\mathcal{M}(\cdot; \nu_t) \neq M^{*,B}(\cdot)) &\leq \mathbb{E}T(\exists u', u'' \in \mathcal{U}, p', p'' \in \mathcal{P} : \mathcal{M}(u'; \nu_t) = p' \cap \mathcal{M}(u''; \nu_t) = p'' \\
 &\quad \cap V(u', p''; \mu) > V(u', p'; \mu) \\
 &\quad \cap V(p'', u'; \mu) > V(p'', u''; \mu) \\
 &\quad \cap (V(u', p''; \nu_t) \leq V(u', p'; \nu_t) \cup V(p'', u'; \nu_t) \leq V(p'', u''; \nu_t)) \\
 &\leq \mathbb{E}T(\exists u', u'' \in \mathcal{U}, p', p'' \in \mathcal{P} : \mathcal{M}(u'; \nu_t) = p' \cap \mathcal{M}(u''; \nu_t) = p'' \\
 &\quad \cap V(u', p''; \mu) > V(u', p'; \mu) \\
 &\quad \cap V(p'', u'; \mu) > V(p'', u''; \mu) \\
 &\quad \cap V(u', p''; \nu_t) \leq V(u', p'; \nu_t)) \\
 &\quad + \mathbb{E}T(\exists u', u'' \in \mathcal{U}, p', p'' \in \mathcal{P} : \mathcal{M}(u'; \nu_t) = p' \cap \mathcal{M}(u''; \nu_t) = p'' \\
 &\quad \cap V(u', p''; \mu) > V(u', p'; \mu) \\
 &\quad \cap V(p'', u'; \mu) > V(p'', u''; \mu) \\
 &\quad \cap V(p'', u'; \nu_t) \leq V(p'', u''; \nu_t)) \\
 &\leq \mathbb{E}T(\exists u \in \mathcal{U}, p', p'' \in \mathcal{P} : \mathcal{M}(u; \nu_t) = p' \\
 &\quad \cap V(u, p''; \mu) > V(u, p'; \mu) \\
 &\quad \cap V(u, p''; \nu_t) \leq V(u, p'; \nu_t)) \\
 &\quad + \mathbb{E}T(\exists u', u'' \in \mathcal{U}, p \in \mathcal{P} : \cap \mathcal{M}(u''; \nu_t) = p \\
 &\quad \cap V(p, u'; \mu) > V(p, u''; \mu) \\
 &\quad \cap V(p, u'; \nu_t) \leq V(p, u''; \nu_t)). \tag{30}
 \end{aligned}$$

Let's split (30) into two terms. Examining the second term,

$$\begin{aligned}
 &\mathbb{E}T(\exists u', u'' \in \mathcal{U}, p \in \mathcal{P} : \mathcal{M}(u''; \nu_t) = p \cap V(p, u'; \mu) > V(p, u''; \mu) \cap V(p, u'; \nu_t) \leq V(p, u''; \nu_t)) \\
 &= \mathbb{E}T(\exists u', u'' \in \mathcal{U}, p \in \mathcal{P} : \mathcal{M}(u''; \nu_t) = p \\
 &\quad \cap \mu(p, u') - c_1 + g(p; \mu) > \mu(p, u'') - c_1 + g(p; \mu) \\
 &\quad \cap \nu_t(p, u') - c_1 + g(p; \nu_t) \leq \nu_t(p, u'') - c_1 + g(p; \nu_t)) \\
 &= \mathbb{E}T(\exists u', u'' \in \mathcal{U}, p \in \mathcal{P} : \mathcal{M}(u''; \nu_t) = p \cap \mu(p, u') > \mu(p, u'') \cap \nu_t(p, u') \leq \nu_t(p, u'')) \\
 &\leq N^2 L \left( \frac{8\sigma^2 \alpha \log(T)}{\Delta_{\min}^2} + \frac{\alpha}{\alpha - 2} \right), \tag{31}
 \end{aligned}$$

where  $\Delta_{\min} = \min_{a, a', a'' \in \mathcal{A}} |\mu(a, a') - \mu(a, a'')|$  and the last inequality follows from Lemma C.4 (for further explanation, see the proof of Lemma C.6).

Turning to the first term of (30),

$$\begin{aligned}
 &\mathbb{E}T(\exists u \in \mathcal{U}, p', p'' \in \mathcal{P} : \mathcal{M}(u; \nu_t) = p' \cap V(u, p''; \mu) > V(u, p'; \mu) \cap V(u, p''; \nu_t) \leq V(u, p'; \nu_t)) \\
 &= \mathbb{E}T(\exists u \in \mathcal{U}, p', p'' \in \mathcal{P} : \mathcal{M}(u; \nu_t) = p' \cap \mu(u, p''; \mu) - c_2 - g(p''; \mu) > \mu(u, p'; \mu) - c_2 - g(p'; \mu) \\
 &\quad \cap \nu_t(u, p'') - c_2 - g(p''; \nu_t) \leq \nu_t(u, p') - c_2 - g(p'; \nu_t)) \\
 &= \mathbb{E}T(\exists u \in \mathcal{U}, p', p'' \in \mathcal{P} : \mathcal{M}(u; \nu_t) = p' \cap \mu(u, p''; \mu) - g(p''; \mu) > \mu(u, p'; \mu) - g(p'; \mu) \\
 &\quad \cap \nu_t(u, p'') - g(p''; \mu) \leq \nu_t(u, p') - g(p'; \mu)) \\
 &\leq N^2 L \left( \frac{8\sigma^2 \alpha \log(T)}{\Delta_{\min}^2} + \frac{\alpha}{\alpha - 2} \right), \tag{32}
 \end{aligned}$$

where the last equality follows from the fact that  $g(\cdot; \mu) = g(\cdot; \nu_t)$ , and the last line follows from Lemma C.4 (for further explanation, see the proof of Lemma C.6).

Combining (28)-(32) yields the regret upper bound as stated in the result.

**Fairness is guaranteed.** That  $\mathcal{M}$  as described in the result also results in fairness (i.e., stability guarantees fairness) follows directly from the analysis under ‘‘Logarithmic optimal regret’’ above. Specifically, since  $\bar{R}(a; \mathcal{M}) = O(\log(T))$  for all  $a \in \mathcal{A}$ , fairness is achieved for all agents by Definition 3.  $\square$