

# Designing Approximately Optimal Search on Matching Platforms

Nicole Immerlica

Microsoft Research, New York, NY, nicimm@gmail.com

Brendan Lucier

Microsoft Research, Cambridge, MA, brlucier@microsoft.com

Vahideh Manshadi

Yale School of Management, New Haven, CT, vahideh.manshadi@yale.edu

Alexander Wei

UC Berkeley, Berkeley, CA, awei@berkeley.edu

We study the design of a decentralized two-sided matching market in which agents' search is guided by the platform. There are finitely many agent types, each with (potentially random) preferences drawn from known type-specific distributions. Equipped with knowledge of these distributions, the platform guides the search process by determining the meeting rate between each pair of types from the two sides. Focusing on symmetric pairwise preferences in a continuum model, we first characterize the unique stationary equilibrium that arises given a feasible set of meeting rates. We then introduce the platform's optimal directed search problem, which involves optimizing meeting rates to maximize equilibrium social welfare. We first show that incentive issues arising from congestion and cannibalization make the design problem fairly intricate. Nonetheless, we develop an efficiently computable search design whose corresponding equilibrium achieves at least  $1/4$  the social welfare of the optimal design. In fact, our construction always recovers at least  $1/4$  the *first-best* social welfare, where agents' incentives are disregarded. Our directed search design is simple and easy-to-implement, as its corresponding bipartite graph consists of disjoint stars. Furthermore, our design implies the platform can substantially limit choice and yet induce an equilibrium with an approximately optimal welfare. Finally, we show that approximation is likely the best we can hope for by establishing that the problem of designing optimal directed search is NP-hard to even approximate beyond a certain constant factor.

*Key words:* directed search, matching platforms, market design, search friction, sharing economy

---

## 1. Introduction

Matching platforms have become prominent facilitators of social and economic connections in recent years: Thirty percent of U.S. adults have used dating platforms to look for a partner (Pew Research Center 2020); and 35% of U.S. workers have engaged in some form of freelance labor in the past year, in part thanks to the growth of online marketplaces for freelance work (Ozimek 2019). Platforms, such as e-Harmony, try to improve search via match recommendations that are based on

learned user characteristics and match compatibility models. While carefully designed models have predictive power, the compatibility of a pair often also involves an *idiosyncratic* component that can only be discovered upon meeting. Furthermore, since matching requires a “coincidence of wants,” the platform cannot compel users to match, even if preferences are known. In such environments, how can matching platforms improve the search process while respecting users’ incentives?

Toward answering this question, we construct a dynamic matching market between two heterogeneous populations, each composed of finitely many types. To be concrete, we take a dating market for heterosexual couples as our base example and thus refer to agents of the two sides as women and men. An agent’s type captures their common features observable to the platform and partially determines their preferences for opposite-side agents.<sup>1</sup> In addition to the type-specific component, each agent’s preference for any opposite-side agent has an idiosyncratic component that is unknown a priori and is only revealed when the pair meets.<sup>2</sup> In our model, agents have cardinal preferences, and each woman-man pair shares a symmetric valuation that is drawn from a distribution corresponding to their types. Symmetric preferences capture the notion of mutual compatibility in a relationship, i.e., any relationship is either win-win or lose-lose. Upon meeting, both agents observe their shared valuation; each then decides whether to match or to wait for another candidate. We consider a continuum model where agents of different types arrive at exogenous rates and leave the market either upon meeting a satisfactory match or unmatched due to a life event which occurs at a given rate (see Section 2.1 and Figure 1). Such an exogenous departure rate captures *search friction* in the sense that an agent can only meet a limited number of candidates before leaving unmatched due to a life event.

With knowledge of preference distributions and exogenous arrival/unmatched departure rates, the platform guides the search process by designing *meeting rates* between pairs of woman-man types. The set of meeting rates for a given type can be viewed as an *assortment* of opposite-side types offered to them over time. Faced with such an assortment, an agent decides on which candidates to accept or reject in order to maximize their long-run utility. We focus on the stationary and symmetric equilibria of the underlying dynamic game: We show for any set of “feasible” meeting rates that respect natural physical constraints (see eqs. (2) and (3)), there exists a unique such equilibrium in which each agent type plays a threshold strategy (see Proposition 3.1). Establishing the uniqueness of the stationary equilibrium relies on the symmetric structure of preferences. In fact, in our constructive proof, we show that iterating the best response map converges after finitely

---

<sup>1</sup> We remark that platforms such as e-Harmony extract these features using surveys; however, we abstract away from such details and assume that these features are directly observable.

<sup>2</sup> Such preference structures have been studied previously in matching literature; see e.g., Ashlagi et al. (2020) and Kanoria and Saban (2021).

many rounds (see Proposition A.4). Uniqueness of equilibrium along with its simple structure enables us to sidestep issues of instability and equilibrium selection which may make the platform’s design problem ill-defined. As such, our structural results may be of independent interest in future work on the design of dynamic matching markets.<sup>3</sup>

We remark that by taking such a design approach to search, our work departs from the prevailing search environment studied in the literature (Burdett and Coles 1997, Shimer and Smith 2000, Adachi 2003, Lauermaun and Nöldeke 2014), which assumes agents meet uniformly at random. In the presence of sufficient differentiation across types, such a “hands-off” approach to search can be arbitrarily sub-optimal: If the platform already knows that sports fans only marry sports fans and outdoor enthusiasts only marry outdoor enthusiasts, it should not waste time by letting outdoor enthusiasts meet sports fans. (For a quantitative example, see the horizontal market presented in Example 3.4.) At the same time, the optimal meeting design can be far more nuanced than simply letting preferred pairs meet each other. This is due to disparities in arrival rates combined with strategic behavior of agents, which can lead to issues of cannibalization and congestion. To alleviate these issues, sometimes the platform may wish to restrict the choices of a type in order to induce a matching outcome with globally higher welfare. To illustrate these behaviors, in Section 3.3, we present an instance of an assortative market with deterministic preferences for which the optimal meeting design is indeed non-assortative (see the vertical market presented in Example 3.5).

The aforementioned behaviors suggest that optimal design can be fairly subtle. In fact, we show that the platform’s optimal directed search problem (formally introduced in Section 3.2) is APX-hard, i.e., NP-hard to approximate beyond a fixed constant factor (see Theorem 4.2). Nonetheless, we develop a solution (i.e., a set of feasible meeting rates) whose equilibrium welfare is at least  $1/4$  of that of the optimal design (see Theorem 4.1). Our design relies on first relaxing the incentive constraints and solving the first best problem. Through an intricate reformulation, we show that the first-best problem can be reduced to a polynomial-time solvable “generalized assignment problem” that admits a forest-structured solution (see Section 4.1). We then show that the assignment problem admits a 2-approximate solution consisting of disjoint “star-shaped” submarkets, i.e., submarkets where one side consists entirely of a single type. Finally, we recompute the first-best for each submarket; utilizing its star structure, we show that the first-best solution can be “modified” to respect incentive constraints with at most a factor of 2 loss in welfare (see Section 4.2, Proposition 4.8, and Figure 2). This holds because for a star-shaped market, there is an alignment of incentives between individual agents and the platform.

---

<sup>3</sup> We highlight that several previous papers considered symmetric preference structures in static settings under other names such as *correlated two-sided* or *acyclic* markets (Abraham et al. 2008, Ackermann et al. 2011).

We complement our theoretical developments with numerical studies of markets where agents’ preferences have both a vertical and a horizontal component (see Section 5.1). Our comparative statics with respect to the strength of each component as well as the intensity of search friction illustrate that, compared to random meeting, our design is particularly effective in markets with strong horizontal components or high search friction (see Figures 3 and 4 in Section 5). Furthermore, our simulation results show that the welfare outcome of our approximation algorithm can be substantially better than its theoretical guarantee.

The simple structure of our proposed solution—a collection of disjoint star-shaped markets—implies that, with careful search design, platforms can substantially limit choice, offer very simple assortments, and yet induce an equilibrium with approximately optimal welfare. By limiting choice, our proposed solution also significantly simplifies the agent’s decision problem and thus reduces the need for complex strategization. Our result is particularly intriguing since it holds without any restrictions on the structure of the symmetric preferences.

### 1.1. Related Work

Our work relates to and contributes to several streams of literature on matching markets.

*Search and Matching.* There is a rich literature in economics that studies decentralized matching models with search frictions and transferable/nontransferable utility. For an informative review, we refer the interested reader to Chade et al. (2017). By and large, this literature focuses on uniformly random meeting among agents and studies the properties of the resulting stationary equilibria. Our framework is partly inspired by the seminal works of Burdett and Coles (1997) and Smith (2006) in this literature. These papers consider matching models with nontransferable utility in which agents of each side belong to a continuum of types. In the work of Burdett and Coles (1997), an agent’s match utility only depends on the type (pizazz) of the other agent, while Smith (2006) considers more general pair-specific utilities (production) that may be symmetric (akin to our preference structure). However, unlike these papers, we do not assume that match utility is deterministic and vertically differentiated.<sup>4</sup> This implies that, in our general model, types are incomparable and assortative matching is not well-defined.<sup>5</sup> (See Example 3.4 for a horizontally differentiated market within our framework.)

Another line of work in this literature studies the relationship between the stationary equilibria of decentralized search under uniformly random meeting and corresponding sets of stable matchings. Under a “cloning” assumption—which keeps the distribution of agents unchanged by asserting that

<sup>4</sup> Chade (2006) extends the model of Burdett and Coles (1997) to incorporate uncertainty in preferences; however, their model differs crucially from ours in that, even upon meeting, only a noisy version of the utility is observed.

<sup>5</sup> After our Example 3.5, which considers a vertically differentiated market, we further discuss the connection between this example and one of the examples presented in Burdett and Coles (1997).

each time two agents match and leave, two single agents identical to them join—Adachi (2003) shows that as search friction vanishes, the set of stationary equilibria converges to the set of stable matchings of a corresponding centralized market. Relaxing such a cloning assumption, and thus allowing for the distribution of agents to be endogenously determined, Lauer mann and Nöldeke (2014) show that stationary equilibria converge to stable matchings if and only if there exists a unique stable matching in a corresponding discrete market.

Following this literature, we also model search friction as a time consuming process of meeting partners and abstract away from incorporating any potential second-stage decision or cost upon meeting (e.g., whether to go on a date with a potential partner). However, we emphasize that we complement this literature by taking a design approach—motivated by online matching platforms—to optimize directed search (i.e., non-uniform and pairwise type-specific meetings) in a model with cardinal and symmetric preferences. In the spirit of designing the search environment, Shimer and Smith (2001) study a model with transferable utility in which search is costly and search intensities can vary across types. (However, agents still meet uniformly at random within the searching population.) They show that a socially optimal solution (i.e., first-best) can arise as an equilibrium of a decentralized search process with a linear tax or subsidy on search intensity.

*Directed Search and Platform Interventions.* Moving beyond random meeting, a series of recent papers has studied different forms of platform design and intervention to facilitate search. Halaburda et al. (2018) focus on the impact of limiting choice and show that when agents’ outside options are heterogeneous, a platform that offers limited choice and charges a higher price can still compete with ones without any choice restriction. A recent work of Kanoria and Saban (2021) shows that limiting the action of agents, e.g., allowing for only one side to propose, and hiding information about the quality of some agents can lead to welfare improvement. We now further discuss some of the key differences between our paper and Kanoria and Saban (2021). First, we remark that the set of actions of agents in Kanoria and Saban (2021) is richer in that, in addition to deciding on whether to match upon meeting, agents decide on whether to request meeting a candidate and whether to inspect that candidate (at a cost). While the model of Kanoria and Saban (2021) also includes exogenous departure rates, their paper focuses on the regime where departure rates approach zero. As a result, each agent has effectively unboundedly many meeting opportunities and the main search friction is inspection cost. In contrast, in our model, each agent has a finite number of meeting opportunities; upon each meeting—after observing the match utility—the agent decides whether to accept the match or to risk departing unmatched before meeting another candidate. (As mentioned above, a majority of papers in the search literature take a similar modeling approach.) Another key difference is the nature of the platform intervention. Kanoria and Saban (2021) investigate blocking one side from proposing and hiding information about quality in the special case of a


vertical market where one side consists of a single type and the other one has two (high/low) types. In contrast, we focus on designing meeting distributions—inspired by platform recommendation systems—for general (not necessarily vertical) markets. Overall, compared to the aforementioned papers, our model and our intervention are more fine-grained, in that both sides can have multiple (potentially incomparable) types and we design pairwise type-specific meeting rates.

The approach of Banerjee et al. (2017) to designing visibility graphs in a two-sided market with different types of buyers and sellers bears some similarity with ours. However, there are several key differences: We consider a matching market without transfers in which the platform’s directed search design impacts agents’ acceptance thresholds on both sides, while Banerjee et al. (2017) focus on a network of buyers and sellers exchanging a single undifferentiated good in which the platform impacts clearing prices by choosing a visibility subgraph.

Motivated by the role of matching platforms in shaping agent’s choice, a sequence of recent papers (Ashlagi et al. 2019b, Aouad and Saban 2020, Ríos et al. 2020) consider the problem of assortment planning in two-sided matching in both static and dynamic settings. In this novel line of work, agents are non-strategic in that their behavior is captured by a choice model such as the Multinomial Logit model. Consequently, agents are oblivious to the action of those on the other side and the platform design. While sharing a similar motivation with this emerging literature, our work complements it by capturing agents’ strategic behavior in response to how the platform’s design guides the search process. In our work, the set of meeting rates for a type can be viewed as an assortment of options that the platform (sequentially) offers to an agent type. The equilibrium response to the collection of such assortments determines the matching and welfare outcome.

*Matching with Incomplete Information.* Outside the framework of search, several papers explore the problem of finding a stable matching where preferences are a priori unknown. Taking a communication complexity approach, Gonczarowski et al. (2019) and Ashlagi et al. (2020) establish bounds on the the amount of communication, measured by the number of bits, needed to find a stable match in markets with private preferences. The recent work of Immorlica et al. (2020) focuses on a setting where learning preferences is costly and show how costly information acquisition impacts an agent’s preference. Furthermore, a few recent papers, such as Liu et al. (2020), use the multi-armed bandit framework to model the process of learning preferences as an online learning problem and develop efficient learning algorithms. Finally, in another direction, Emamjomeh-Zadeh et al. (2020) analyze an iterative query process for learning a stable matching under general preferences. A closely related line of work focuses on specific matching algorithms such as the Men-Proposing Deferred Acceptance Algorithm, and examines agents’ incentives to manipulate preferences (see for example Roth (1982), Coles and Shorrer (2014), and Immorlica and Mahdian (2015)).

**Figure 1** This schematic depicts the dynamics for a single type  $\theta$ . Each dark gray box represents a type, with the size corresponding to the type’s population mass. The light gray connections between boxes correspond to how agents meet under directed search. The colored lines depict inflow and outflow for type  $\theta$ , with width corresponding to flow rate: Blue for arrivals, green for matching, and red for leaving unmatched. In stationary equilibrium, the inflow should balance the outflow (as is depicted).



*Matching in Dynamic Environments.* A growing number of papers study the design of dynamic matching markets motivated by various applications ranging from school choice (Feigenbaum et al. 2020) to ride-sharing (Nikzad 2017) to public housing (Afeche et al. 2021). In this body of literature, several papers focus on the timing (or frequency) of matching in a centralized market and the trade-off between market thickness and waiting time (Ashlagi et al. 2013, 2019a, Akbarpour et al. 2020). In a closely related direction, papers such as Doval and Szentes (2019) and Arnosti and Shi (2020) study the setting where strategic agents have heterogeneous preferences and thus face a trade-off between matching with a less preferred choice (including their outside option) or waiting longer. Recently, several papers examine the role of information design on managing congestion in dynamic allocation problems (Anunrojwong et al. 2021, Ashlagi et al. 2021, Che and Tercieux 2021). Our work complements the aforementioned papers by studying the design of a dynamic matching market in which strategic agents with heterogeneous preferences repeatedly meet potential matches recommended by the platform.

## 2. Model

In this section, we introduce a search model in a dynamic, two-sided matching market, where the search process is facilitated by recommendations from a centralized platform. A graphical depiction of our model, focused on the dynamics around a single type, is shown in Figure 1.

## 2.1. Agents, Types, and Dynamics

We model our matching market as a flow economy, where the market consists of a continuum of agents of infinitesimal mass. Agents on the two sides are differentiated into finitely many types belonging to the sets  $\mathcal{M}$  and  $\mathcal{W}$ , respectively. Agents' types determine prior distributions over their preferences (see Section 2.2). We use  $\Theta = \mathcal{M} \cup \mathcal{W}$  to denote the combined type space of all agents. For each type  $\theta \in \Theta$ , we denote the population mass of type  $\theta$  agents present on the platform by  $\eta_\theta$ . Throughout, we will use the notational convention that  $\theta'$  refers to a generic type on the opposite side of the market from type  $\theta$ .

Our matching market takes place in continuous time. To track change in the population masses  $\eta_\theta$  over time, we use *flow rates*—the mass of agents arriving or departing per unit time.<sup>6</sup> We assume entry into the market is fixed exogenously, with type  $\theta$  agents entering at a constant flow rate of  $\alpha_\theta$  (i.e., an  $\alpha_\theta$  mass of type  $\theta$  agents enters per unit time). This corresponds to the blue inflow in Figure 1. On the other hand, we let departure from the market be partially determined endogenously, with all agents eventually either leaving with a match or experiencing an exogenous “life event” that causes them to leave unmatched. More specifically, the two ways an agent can depart are as follows:

**Matching.** Agents leave the market when they enter into a mutually agreed upon match with another agent (see the green outflow in Figure 1). We define  $\xi_\theta$  so that  $\xi_\theta \eta_\theta$  is the flow rate at which type  $\theta$  agents do so. That is, each *individual* of type  $\theta$  (assuming symmetry between agents of type  $\theta$ ) leaves matched with probability  $\xi_\theta dt$  during the infinitesimally small time interval  $[t, t + dt)$ . The value of  $\xi_\theta$  is determined both by how the platform recommends potential matches and who the agents themselves choose to match with (see Lemma 3.2).<sup>7</sup>

**Life event.** Agents also leave unmatched when they experience a “life event” (see the red outflow in Figure 1). We assume life events occur randomly for each individual at a constant rate of  $\delta dt$ . That is, each individual experiences a life event with probability  $\delta dt$  during each infinitesimally small time interval  $[t, t + dt)$ . (Equivalently, the time an agent spends in the market is drawn (unknown to them) from an exponential distribution of rate  $\delta$ .) At the individual level, the possibility of leaving unmatched acts as a discount factor and ensures that agents do not search forever. At the platform level, individuals leave at a total flow rate of  $\delta \eta_\theta$ , which ensures that the mass of unmatched agents does not grow unboundedly.

<sup>6</sup> Although  $\eta_\theta$  can in general vary over time, we suppress the time index on  $\eta_\theta$  for notational compactness and because our eventual focus will be on stationary equilibria, in which case  $\eta_\theta$  will be constant over time.

<sup>7</sup> Note that  $\xi_\theta$  can also vary over time; we suppress the time index as we did for  $\eta_\theta$ .



Our focus will be on the stationary equilibria of this market, where the inflow of arriving agents balances the outflow of departing agents. We say that *stationarity* holds if

$$\alpha_\theta = (\delta + \xi_\theta)\eta_\theta \tag{1}$$

for all  $\theta \in \Theta$ . Note that  $\delta + \xi_\theta$  is the total rate at which *individual* type  $\theta$  agents exit the market.

## 2.2. Preferences and Utilities

We assume that agents have cardinal preferences (i.e., a utility associated to matching with each potential partner) and normalize the value of leaving unmatched to 0. In our model, the utility values of specific potential partners are *unknown* ex ante. Instead, pairs of agents learn their valuations for each other when they meet while searching (see Section 2.3), with their utilities being drawn from a prior distribution given by their types.<sup>8</sup> These priors represent the fact that agents may have idiosyncratic preferences—beauty being in the eye of the beholder—that are not captured by their type. We assume that these priors are common knowledge among the platform and the agents.

More formally, the utilities that two agents of types  $m \in \mathcal{M}$  and  $w \in \mathcal{W}$  derive if they match are drawn from some joint distribution  $\mathcal{D}_{mw}$ . Because we use  $\mathcal{D}_{mw}$  to capture idiosyncrasies in preferences, we assume that utilities are drawn independently from  $\mathcal{D}_{mw}$  for each such pair of agents. To model mutual compatibility, we also assume agents have *symmetric valuations*, meaning that if two agents match, then each agent derives the same utility  $u \in \mathbb{R}$  from the match.

As mentioned in the introduction, an agent’s type captures their common features. In our abstract model, we assume that the platform directly observes these features and thus the agent’s type. While some features in practice may be self-reported by agents and thus present opportunities for misreporting, we argue that this is not a major concern in our setting and motivating applications. In practice, repeated meetings make verification easier: For example, if one misreports their height, then that would be revealed after meeting a potential partner, who could report back to the platform. At the same time, untruthful reporting may result in severe punishments for an agent, such as being permanently banned from the platform.<sup>9</sup> On a more abstract level, in many settings of interest, e.g., horizontally differentiated markets, it is unclear whether misreporting even helps an agent.

Under the symmetric valuations assumption, each joint distribution  $\mathcal{D}_{mw}$  can be thought of as a distribution over  $\mathbb{R}$  with cumulative distribution function  $F_{mw}$ . To avoid a technical discussion of

---

<sup>8</sup> We assume that agents do not meet twice, since for any agent, there are uncountably many agents of each type that they could meet.

<sup>9</sup> For example, Hinge’s Terms of Service require each user to not “*Misrepresent your identity, age, current or previous positions, qualifications, or affiliations with a person or entity.*” Violating these terms results in being banned from the platform.

tiebreaking<sup>10</sup> we make the mild assumption that  $\mathcal{D}_{mw}$  is a continuous distribution for all  $(m, w)$ . Then  $F_{mw}$  is absolutely continuous and  $F_{mw}(\tau) = \Pr(u \leq \tau)$  for  $u \sim \mathcal{D}_{mw}$ . Finally, to make our notation for distributions symmetric, we define  $\mathcal{D}_{wm} := \mathcal{D}_{mw}$  and  $F_{wm} := F_{mw}$  for all  $(m, w)$ .

### 2.3. Directed Search

A key feature of our model is that search is mediated by a platform. At a high level, the platform directs search by picking for each agent type  $\theta$  the rates at which they “meet” other agent types while searching. For each agent, search consists of a sequence of meetings with potential partners until either a match is formed or they leave unmatched. When two agents meet, each must decide whether to accept or reject the match. If both parties accept (i.e., there is a double coincidence of wants), then they match with each other and leave the market. Otherwise, both agents remain in the market and continue their respective searches.

At the individual level, the frequency at which an agent encounters other agents of any given type follows a Poisson process with rate given by the platform’s directed search policy. Upon two agents of types  $\theta$  and  $\theta'$  meeting, their valuation  $u$  for the match is drawn from  $\mathcal{D}_{\theta\theta'}$  and revealed to them. Each agent then decides whether to accept or reject the match, with the match forming if and only if both accept. If the agents match, each leaves the market with utility  $u$ ; otherwise, each continues their search. We assume agents are rate-limited and normalize time so that agents can meet at most one candidate per unit time on average. That is, we require the total frequency of meetings for any agent to follow a Poisson process with at most unit rate—this is our *capacity constraint*.

At the platform level, directed search can be thought of as some mass  $f_{mw}$  of type  $m$  and type  $w$  agents, for each pair of types  $(m, w)$ , meeting each other per unit time. In particular, *equal* masses of type  $m$  and type  $w$  agents should meet each other over any time interval. Note that this property must be satisfied if we were to explicitly pair agents in a discrete model. While our model is continuous, we can still imagine the platform “pairing” equal masses of randomly sampled agents for meetings at each instant of time. By requiring that agent-level Poisson processes be consistent with this fact, we obtain our *flow balancing* constraint.

We now state the directed search model sketched above in more formal terms. For any  $\theta \in \Theta$ , define an *assortment* to be a collection of rates  $\lambda_\theta(\theta') \geq 0$  for types  $\theta'$  on the opposite side such that type  $\theta$  agents meet type  $\theta'$  agents following a Poisson process with rate  $\lambda_\theta(\theta')$ . The platform, as part of its design, chooses an assortment  $\lambda_\theta$  for each type  $\theta \in \Theta$ .

A set of assortments  $\{\lambda_\theta\}_{\theta \in \Theta}$  is *feasible* if it satisfies the following two sets of constraints:

<sup>10</sup> While all of our arguments work with discrete distributions, one has to take care to define tiebreaking properly. On the other hand, note that any non-continuous distribution can be approximated by a continuous distribution by adding a small amount of (bounded) noise.

**Capacity.** The total rate of meetings for type  $\theta$  agents is bounded by 1 for all types  $\theta \in \Theta$ :

$$\sum_{\theta'} \lambda_{\theta}(\theta') \leq 1. \quad (2)$$

(Note that  $\sum_{\theta'} \lambda_{\theta}(\theta')$  is the total rate because merging Poisson processes sums their rates.)

**Flow balance.** For all pairs  $(m, w) \in \mathcal{M} \times \mathcal{W}$ , the mass of type  $w$  agents met by type  $m$  agents per unit time equals the mass of type  $m$  agents met by type  $w$  agents per unit time. By an exact law of large numbers (Sun 2006, Duffie et al. 2018)<sup>11</sup>, we can write this condition as

$$f_{mw} := \eta_m \lambda_m(w) = \eta_w \lambda_w(m). \quad (3)$$

That is, the mass  $f_{mw}$  of type  $m$  and  $w$  agents that meet per unit time is well-defined.

We assume the platform can direct search according to any feasible set of assortments.

To motivate the Poisson assumption, we think of the platform as implementing a feasible set of assortments by randomly sampling a subset of each type to meet some other type at each instant in time. With uniform sampling, individual agents will experience meetings at a Poisson rate. Note that by assuming uniform sampling, we implicitly restrict the platform to designs that are anonymous and independent of agents' histories.

### 3. The Platform's Design Problem

In this section, we introduce the platform's design problem, namely: *How should the platform suggest meetings in directed search to maximize social welfare in equilibrium?* In order to formally define the problem, we first describe the agents' strategic decision-making problem and briefly characterize the equilibria that arise.

#### 3.1. Strategies and Equilibrium Play

Any set  $\{\lambda_{\theta}\}_{\theta \in \Theta}$  of assortments defines a game for the agents, in which agents strategically decide which potential partners to accept and reject in their meetings. As our focus will be on stationary equilibria, we assume agents play *time-* and *history-independent* strategies. We further assume that strategy profiles are *symmetric* within each type. We can thus write the strategy of type  $\theta$  agents as a function  $\sigma_{\theta}(\theta', u)$  taking values in  $[0, 1]$ , specifying the probability with which a type  $\theta$  agent accepts when meeting a type  $\theta'$  agent whom they value at utility  $u$ .

A property of our setup is that no individual agent can directly influence the behavior of other agents, as agents are only affected by the action profiles of types *in aggregate*. Thus, rather than modeling agents as playing a full-fledged dynamic game, we can think of each agent as facing a

---

<sup>11</sup> As in the related literature on search and matching, we formally require a continuous-time exact law of large numbers for random matching; such a result has only recently been developed rigorously for discrete-time settings by Duffie et al. (2018).

Markov decision process (MDP) derived from their assortment  $\lambda_\theta$  and strategies  $\sigma_{\theta'}$  of the agents on the opposite side.<sup>12</sup> Then, a strategy  $\sigma_\theta$  is a best response for type  $\theta$  agents if and only if it is an optimal policy for this MDP. It follows that a strategy profile  $\{\sigma_\theta\}_{\theta \in \Theta}$  is a Nash equilibrium if and only if the strategy of each type is an optimal policy for the MDP given by their assortment and opposite side's strategies.

In Section A.1, we formally describe the agents' MDPs and show that any Nash equilibrium in “non-dominated” strategies is equivalent to one in *threshold strategies*, i.e., strategies  $\sigma_\theta$  for which there exists a threshold  $\tau_\theta$  such that  $\sigma_\theta(\theta', u) = 1$  if  $u \geq \tau_\theta$  and  $\sigma_\theta(\theta', u) = 0$  if  $u < \tau_\theta$ .<sup>13</sup> In fact, each threshold  $\tau_\theta$  in this equilibrium equals the expected utility  $\bar{u}_\theta$  of type  $\theta$  agents upon entry. (Note that  $\bar{u}_\theta$  depends on type  $\theta$ 's strategy, the other agents' strategies, as well as the platform's assortments—see Lemma 3.3.) Furthermore, we show in Section A.3 that the equilibrium prediction of our model is unique in the following sense: For any set of assortments, our model makes a unique prediction of the distribution of matches that are realized. We record these observations in the following proposition (proven in Section A.3):

**Proposition 3.1** (Structure of Nash Equilibria). *Given any fixed set of assortments  $\{\lambda_\theta\}_{\theta \in \Theta}$ :*

1. *There exists a unique strategy profile where each type plays a threshold strategy with threshold equal to that type's expected utility. Moreover, this strategy profile is a Nash equilibrium.*
2. *Any Nash equilibrium in non-dominated strategies produces the same distribution (in both types and utilities) of realized matches as the preceding Nash equilibrium in threshold strategies.*

In light of Proposition 3.1, it suffices to focus our attention on equilibria in threshold strategies. To facilitate our discussion of such equilibria, we state expressions for the rate  $\xi_\theta$  at which agents match (in Lemma 3.2) and their expected utility  $\bar{u}_\theta$  when all agents play threshold strategies (in Lemma 3.3).

**Lemma 3.2** (Matching Rate). *Suppose each type  $\theta \in \Theta$  plays a threshold strategy with threshold  $\tau_\theta$ . Then the rate  $\xi_\theta$  at which individual type  $\theta$  agents match is*

$$\xi_\theta = \sum_{\theta'} \left( \lambda_\theta(\theta') \int_{\max(\tau_\theta, \tau_{\theta'})}^{\infty} dF_{\theta\theta'} \right). \quad (4)$$

*Proof.* The rate at which an agent of type  $\theta$  matches is the sum of the rates at which they meet each type  $\theta'$  weighted by the probability such a meeting results in a match. The latter probability is  $1 - F_{\theta\theta'}(\max(\tau_\theta, \tau_{\theta'})) = \int_{\max(\tau_\theta, \tau_{\theta'})}^{\infty} dF_{\theta\theta'}$  in terms of the agents' thresholds  $\tau_\theta$  and  $\tau_{\theta'}$ .  $\square$

<sup>12</sup>Note that agents are not directly affected by the actions of other agents on the same side.

<sup>13</sup>In our analysis, we eliminate dominated strategies that reject matches worth more than one's expected continuation utility. Performing such a pruning is necessary to rule out degenerate equilibria (e.g., the one where agents reject all of their potential matches because they do not expect anyone to ever accept).

**Lemma 3.3** (Expected Utility). *Suppose each type  $\theta$  plays a threshold strategy with threshold  $\tau_\theta$ . Then the expected payoff  $\bar{u}_\theta$  of type  $\theta$  agents is*

$$\bar{u}_\theta = \frac{\sum_{\theta'} \left( \lambda_{\theta'}(\theta) \int_{\max(\tau_\theta, \tau_{\theta'})}^{\infty} u dF_{\theta\theta'} \right)}{\delta + \sum_{\theta'} \left( \lambda_{\theta'}(\theta) \int_{\max(\tau_\theta, \tau_{\theta'})}^{\infty} dF_{\theta\theta'} \right)}. \quad (5)$$

(Note that by Lemma 3.2, the denominator can also be written as  $\delta + \xi_\theta$ .)

While the proof of Lemma 3.3 depends on the formal definition of the MDP, the intuition for the formula is simple: For any type  $\theta$  agent, the memorylessness of the MDP implies their expected utility equals their expected utility conditioned on leaving at any particular instant. Conditioning on such an event, we see that with probability proportional to  $\delta$ , they left due to a life event, and with probability proportional to  $\lambda_{\theta'}(\theta) \int_{\max(\tau_\theta, \tau_{\theta'})}^{\infty} F_{\theta\theta'}$ , they left because they matched with an agent of type  $\theta'$ . In equation (5), we simply express the agent's expected utility as the sum of their expected utilities in each of these scenarios weighted by probability of occurrence.

### 3.2. Optimal Directed Search

The platform's design problem is to design the search process via assortments in order to induce a stationary equilibrium that is (approximately) optimal in terms of social welfare. Thus, in addition to designing the assortments, the platform should ensure that those assortments are sustainable in equilibrium. When analyzing the design problem, we assume that the platform has access to the departure rate  $\delta$ , the arrival rates  $\alpha_\theta$  for all  $\theta \in \Theta$ , and the preference distributions  $\mathcal{D}_{mw}$  for all  $(m, w) \in \mathcal{M} \times \mathcal{W}$ .

More formally, we say that a set of assortments  $\{\lambda_\theta\}_{\theta \in \Theta}$  induces a *stationary equilibrium* if, for the equilibrium strategy profile  $\{\sigma_\theta\}_{\theta \in \Theta}$ , there exist population masses  $\{\eta_\theta\}_{\theta \in \Theta}$  such that:

1. Stationarity (as defined in eq. (1)) holds for all types  $\theta \in \Theta$ .
2. The assortments are feasible with respect to the population masses (i.e., (2) and (3) hold).

By Proposition 3.1, a set of assortments  $\{\lambda_\theta\}_{\theta \in \Theta}$  implies a unique Nash equilibrium in threshold strategies. Furthermore, these threshold strategies determine  $\xi_\theta$  for each  $\theta$  by Lemma 3.2. Given  $\{\xi_\theta\}_{\theta \in \Theta}$ , the stationarity condition then determines  $\eta_\theta$  for each  $\theta$ . Therefore, any choice of assortments  $\{\lambda_\theta\}_{\theta \in \Theta}$  induces at most one stationary equilibrium, making the platform's optimization problem over the assortments well-defined.<sup>14</sup>

Recall that the platform's optimization objective is social welfare. For stationary equilibria, social welfare is defined as the total flow rate of utility realized by agents matching with each other. Note that the symmetric valuations assumption implies each of the two sides contributes exactly half of

<sup>14</sup>Since our model uniquely predicts play given assortments, stationary equilibria are also self-sustaining once established.

the total welfare. The platform’s optimization problem is thus to find the assortments  $\{\lambda_\theta\}_{\theta \in \Theta}$  that induce a stationary equilibrium with maximum welfare.

Our setup naturally lets us cast the platform’s design problem of finding a welfare-maximizing policy for directed search as a computational one: Given inputs  $\delta$ ,  $\{\alpha_\theta\}_{\theta \in \Theta}$ , and  $\{\mathcal{D}_{mw}\}_{(m,w) \in \mathcal{M} \times \mathcal{W}}$ , *compute* the assortments  $\{\lambda_\theta\}_{\theta \in \Theta}$  that induce a stationary equilibrium with maximum welfare. We name this problem `OptimalDirectedSearch`.<sup>15</sup>

We remark that we do not dwell on dynamical issues beyond steady state (e.g., convergence to equilibrium from “cold starts”). Instead, we assume that the platform has the power to manipulate entry into the market at time 0 and can ensure that the requisite populations of agents are present for the desired equilibrium to sustain indefinitely. This motivates our objective for the platform of finding a welfare-maximizing stationary equilibrium.

**3.2.1. Optimal Directed Search as an Optimization Problem** Stationarity, feasibility, and agents playing best responses can all be encoded (albeit indirectly) as constraints on the assortments. We can thus write `OptimalDirectedSearch` as the following constrained optimization problem:

$$\max_{\lambda_\theta, \tau_\theta, \xi_\theta, \eta_\theta} \sum_{m \in \mathcal{M}} \sum_{w \in \mathcal{W}} \left( \eta_m \lambda_m(w) \int_{\max(\tau_m, \tau_w)} u dF_{mw} \right) + \sum_{w \in \mathcal{W}} \sum_{m \in \mathcal{M}} \left( \eta_w \lambda_w(m) \int_{\max(\tau_m, \tau_w)} u dF_{mw} \right) \quad (6)$$

$$\text{such that } \alpha_\theta = (\delta + \xi_\theta) \eta_\theta \quad \forall \theta \in \Theta \quad (6a)$$

$$\eta_m \lambda_m(w) = \eta_w \lambda_w(m) \quad \forall (m, w) \in \mathcal{M} \times \mathcal{W} \quad (6b)$$

$$1 \geq \sum_{\theta'} \lambda_\theta(\theta') \quad \forall \theta \in \Theta \quad (6c)$$

$$\xi_\theta = \sum_{\theta'} \left( \lambda_\theta(\theta') \int_{\max(\tau_\theta, \tau_{\theta'})}^\infty dF_{\theta\theta'} \right) \quad \forall \theta \in \Theta \quad (6d)$$

$$\tau_\theta = \frac{1}{\delta + \xi_\theta} \sum_{\theta'} \left( \lambda_\theta(\theta') \int_{\max(\tau_\theta, \tau_{\theta'})}^\infty u dF_{\theta\theta'} \right) \quad \forall \theta \in \Theta \quad (6e)$$

$$\lambda_m(w), \lambda_w(m) \geq 0 \quad \forall (m, w) \in \mathcal{M} \times \mathcal{W}. \quad (6f)$$

Here, constraint (6a) is the stationarity constraint (1); constraints (6b) and (6c) are the feasibility constraints of flow balance (3) and capacity (2); constraint (6d) is formula (4) for individuals’ rate of matching; constraint (6e) states that each type’s threshold is a fixed point of their MDP payoff function (5); and constraint (6f) simply requires that the meetings rates be non-negative.

<sup>15</sup> To formalize our computational problem while avoiding details relating numerical computation, we assume that the platform’s knowledge of the preference distributions  $\mathcal{D}_{mw}$  comes in the form of oracle access to the CDFs  $F_{mw}$  and the tail expectations  $\int_\tau^\infty u dF_{mw}$ . We further assume that the platform can solve for each agent’s best response optimally. This is a mild assumption because the corresponding optimization problem boils down to a ternary search on a unimodal function of the aforementioned oracle’s outputs (see Lemma 4.5 and its generalization Lemma B.1 in Section A).

The validity of constraint (6e) follows from Proposition 3.1: Our model’s unique prediction for which matches are realized occurs when agents all play threshold strategies with threshold equal to their expected utility. Moreover, the profile of thresholds satisfying this property is uniquely determined by the assortments. Hence the social welfare of an assortment is given by the flow of utility realized when all agents play such threshold strategies.

We note that while the optimization program (6) is written as an optimization problem over four sets of variables  $\{\lambda_\theta\}_{\theta \in \Theta}$ ,  $\{\tau_\theta\}_{\theta \in \Theta}$ ,  $\{\xi_\theta\}_{\theta \in \Theta}$ , and  $\{\eta_\theta\}_{\theta \in \Theta}$ , it is really an optimization problem over the assortments  $\{\lambda_\theta\}_{\theta \in \Theta}$ : Agents’ thresholds  $\{\tau_\theta\}_{\theta \in \Theta}$  and matching rates  $\{\xi_\theta\}_{\theta \in \Theta}$  are determined entirely by their assortments, and the stationary masses  $\{\eta_\theta\}_{\theta \in \Theta}$  are in turn determined entirely by the matching rates.

### 3.3. Illustrative Examples: Directed Search in Horizontal and Vertical Markets

With our setup of the model and the platform’s design problem now complete, we present two simple examples with fundamentally different market structures: a horizontal market and a vertical one. In these special cases, we are able to optimally solve `OptimalDirectedSearch` and thus gain insight into the structure and value of the optimal design. Through the former example, we show that platform-directed search can provide significant welfare gains relative to random meeting. Through the latter, we illustrate some subtleties involving congestion and cannibalization that arise (even in this very restricted setting) when solving `OptimalDirectedSearch`.

To motivate the utility of platform-directed search, we describe a market involving highly heterogeneous preferences (i.e., horizontal differentiation) where directed search leads to significantly greater efficiency than random meeting.

**Example 3.4** (Horizontal market). We consider a market where the two sides are  $\mathcal{M} = \{m_1, \dots, m_n\}$  and  $\mathcal{W} = \{w_1, \dots, w_n\}$ , with types  $m_i$  and  $w_j$  achieving positive utility from matching only if  $i = j$ . Suppose also that the types are symmetric, with each type having the same arrival rate  $\alpha_\theta = 1$  and each pair of compatible types  $(m_i, w_i)$  having the same preference distribution  $F_{m_i w_i} = F$ . Finally, let  $\delta > 0$  be arbitrary.

We compare the welfares of two stationary equilibria—one under random meeting and another under optimal directed search. Consider the symmetric equilibrium where agents meet randomly. By symmetry, only a  $1/n$  fraction of each agent’s meetings are with agents of the compatible type. On the other hand, with the optimal directed search, the platform sets assortments so that  $\lambda_{m_i}(w_i) = \lambda_{w_i}(m_i) = 1$  for all  $i$ .<sup>16</sup> Then, agents meet other agents of the same type with probability 1. To compare these two equilibria, observe that the random meeting equilibrium is equivalent to the optimal directed search equilibrium with search friction  $n\delta$ —under random meeting, each agent

<sup>16</sup> This policy is optimal because its welfare matches the that of the first-best outcome (defined in Section 4.1).

has a  $n$  times higher probability of leaving unmatched between meetings with compatible agents. In the limit of large  $\delta$ , where agents are likely to accept on their first meeting (or not match at all), this leads to a multiplicative gap approaching  $n$  between the welfares of random meeting and optimal directed search.

This example demonstrates that in a market with horizontal preferences, there can be a significant amount of wasteful meeting if agents proceed with random search: In our example, under random meeting, only one in  $n$  meetings involves pairs that are compatible with each other. On the other hand, with directed search, the platform could eliminate such wasteful meeting entirely. While admittedly stylized, this example nonetheless highlights the usefulness of directed search in horizontally differentiated markets.

In our second example, of a vertically structured market, agents on each side of the market have a common preference ordering over the types on the opposite side. In such a market, one might expect that the optimal directed search scheme results in a positively assortative matching when agents' valuations are supermodular. However, this turns out not to always be the case. Because agents arrive at disparate rates, matching assortatively can lead to inefficiency due to congestion at the top of the market, and due to “high type” agents cannibalizing demand for “low type” agents. While such phenomena make finding the exact optimal assortments challenging, we will show in Section 4 that we can nonetheless find a set of assortments that are approximately optimal.

**Example 3.5** (A Vertical Market). Consider a market with strictly vertical preferences, where the agents on each side belong to either a high type  $H$  or a low type  $L$ . That is, we let  $\mathcal{M} = \{m_H, m_L\}$  and  $\mathcal{W} = \{w_H, w_L\}$ . Suppose the utility distributions  $F_{m_H w_H}$ ,  $F_{m_L w_H}$ ,  $F_{m_H w_L}$ , and  $F_{m_L w_L}$  are point masses at  $u(1 + 2\varepsilon)$ ,  $u$ ,  $1 + \varepsilon$ , and  $1$ , respectively, for  $u \gg 1$  and  $\varepsilon \ll 1$ .<sup>17</sup> Thus, agents always prefer matching with high-type agents over low-type agents. Suppose further that the arrival rates are such that  $\alpha_{m_H} = \alpha_{w_L} = 1$  and  $\alpha_{w_H} = \alpha_{m_L} = 1/u$ . Finally, we consider this example in the regime where  $\delta$  is very small (i.e., suppose we are in a nearly frictionless market).

To analyze the possible stationary equilibria of this market, we first consider what happens in the equilibria where high-types only match with high-types and low-types only match with low-types. The maximum possible welfare in such an equilibrium is bounded above by  $2((1 + 2\varepsilon) + \alpha_{m_L})/(1 + \delta) \approx 2$ , since the welfare generated by the high-types is at most  $2 \cdot \alpha_{w_H} \cdot u(1 + 2\varepsilon)/(1 + \delta)$  and the utility generated by the low-types is nearly negligible at  $2 \cdot \alpha_{m_L} \cdot 1/(1 + \delta)$ . In particular, matching for both high- and low-types is bottlenecked by the low arrival rate of the less common type. One can

<sup>17</sup> An astute reader may notice that, technically, our preference distributions do not satisfy our continuity assumption. We present our example this way for simplicity's sake—the same example can be made to work with continuous distributions by adding some small amount of Gaussian noise to each utility value.



think of these equilibria as being congested at both levels of the market: Due to the disparities in arrival rates, there are many type  $m_H$  agents hoping to match with only a few type  $w_H$  agents and many type  $w_L$  agents hoping to match with only a few type  $m_L$  agents. As a result, large numbers of type  $m_H$  agents and type  $w_L$  agents go unmatched.

Next we observe that such equilibria—where high-types only match with high-types and low-types only match with low-types—arise if the expected utility  $\bar{u}_{m_H}$  of type  $m_H$  is greater than  $1 + \varepsilon$ . Notably, type  $m_H$  agents *are not willing* to match with type  $w_L$  agents due to the possibility of matching with a type  $w_H$  agent: If  $\bar{u}_{m_H} > 1 + \varepsilon$ , then Proposition 3.1 tells us that type  $m_H$  agents will not match with type  $w_L$  agents, since type  $m_H$  agents would only obtain  $1 + \varepsilon$  utility from doing so. From the platform’s perspective, the presence of  $w_H$  agents in type  $m_H$ ’s assortment cannibalizes demand for type  $w_L$  agents. Likewise, the possibility of matching with type  $m_H$  makes type  $w_H$  reject type  $m_L$  agents—the expected utility of type  $w_H$  agents is at least  $(1 + \varepsilon) \cdot \alpha_{m_H} / \alpha_{w_H} > u$  from matching with type  $m_H$ .

On the other hand, we can obtain a more efficient equilibrium if the expected utility of type  $m_H$  is slightly lower, e.g., if they only match with type  $w_L$ . Consider the assortment where type  $m_H$  and  $w_L$  agents only meet each other and type  $w_H$  agents and  $m_L$  agents only meet each other. All agents meet at maximal rates; furthermore, all meetings result in matches because no type has other options. As a result, the social welfare in equilibrium is  $2(\alpha_{m_H} \cdot (1 + \varepsilon) + \alpha_{m_L} \cdot u) / (1 + \delta) = 2((1 + \varepsilon) + 1) / (1 + \delta) \approx 4$ . That is, the social welfare of this non-assortative “diagonal” set of assortments (which happens to be optimal) is nearly twice that of any equilibrium where high-types only match with high-types and low-types only match with low-types.

We conclude our discussion of the above market with two remarks. First, the cannibalization effect (and its welfare implications) is in the same spirit as the *sorting externality* discussed in Burdett and Coles (1997) using a high/low type vertical market, but with a different preference structure. In that setting, the preference of an agent only depends on the type of the other side, and agents meet each other uniformly at random. Depending on the belief that type  $H$  would match with a type  $L$  or not (and its corresponding match threshold), two different equilibria arise. Interestingly, the assortative equilibrium has the lower welfare compared to the equilibrium under which type  $H$  matches both types. Next, we note these phenomena arise in part due to the fact that the platform design is anonymous and thus it cannot discriminate between agents beyond their type. Indeed, at the opposite extreme, the platform could simply force agents to match with a certain type, e.g., by having a subset of type  $m_H$  agents only meet type  $w_L$  agents and another subset of type  $m_H$  agents only meet type  $w_H$  agents. However, we stress that such splitting of types has two key drawbacks: (i) if agents are able to rejoin the platform under a different identity,

then such separation is difficult to enforce; (ii) such discriminatory treatment may be unpalatable to users from a fairness perspective. We also show in Section 4 that such an assumption is not restrictive—our solution is constant-factor competitive against the *first-best* platform policy.

## 4. Approximation Algorithms

In this section, we show how the platform can construct in polynomial time a set of assortments  $\{\lambda_\theta\}_{\theta \in \Theta}$  such that its induced equilibrium obtains a 4-approximation to the optimal social welfare. Furthermore, the equilibrium resulting from our computed assortments has a simple and appealing structure: It consists of disjoint submarkets such that each submarket has only a single type on one of the sides. We also establish hardness of approximation for the platform’s optimization problem, meaning that we cannot do better than a constant-factor approximation unless  $P = NP$ .

We actually show a slightly stronger result: The resulting equilibrium obtains a 4-approximation to the optimal welfare that could be obtained in the first-best matching economy, where rather than agents choosing to accept or reject, the platform gets to decide on behalf of the agents. (In this first-best economy, the platform gets to choose the assortments as well.) Our theorem can thus also be interpreted as a price-of-anarchy-style result, that agents’ self-interested behavior causes welfare to degrade by at most a factor of 4 relative to the first-best optimum.

**Theorem 4.1** (Computationally Efficient Approximation). *There is a polynomial-time 4-approximation algorithm for OptimalDirectedSearch. (In fact, this algorithm gets a 4-approximation to the platform’s first-best optimal welfare.)*

Our construction of these assortments and the resulting equilibrium proceeds as follows:

**Solving for the platform’s first-best.** We start by considering the platform’s optimization problem for the first-best matching economy, where the platform gets to plan both the assortments and which matches the agents accept. That is, we relax the agents’ incentive constraint and allow the platform to optimize over thresholds  $\tau_{mw}$  between each pair  $(m, w)$  of types. This relaxation removes the main barrier—namely, incentives—to solving the optimization problem exactly. Accordingly, we show that this relaxed problem can be reduced to a generalization of the classical assignment problem and thus be efficiently solved.

**Approximating via “star-shaped” markets.** We then show that an optimal solution to this “generalized” assignment problem can be converted into a stationary equilibrium (where we now take agents’ incentives back into account) with at most a factor 4 loss in welfare. The key ideas for this step are:

- Showing that the optimal first-best solution can be 2-approximated by another consisting of “star-shaped” submarkets (which have a side where all the agents are of the same

type). For this step, we draw on some classical observations about the structure of optimal solutions from previous works on generalized assignment problems (Lenstra et al. 1990, Chakrabarty and Goel 2010, Banerjee et al. 2017).

- Showing that for markets where all the agents are of the same type on one side, the platform can obtain a 2-approximation to the optimal welfare of the first-best matching economy. This step relies on the observation that in such markets, there is an alignment of incentives between individual agents and the platform.

A constant-factor approximation algorithm is likely the best we can hope for, computationally: We show in Section C that approximating the platform’s optimization problem to a factor better than  $\frac{24}{23}$  is NP-hard.

**Theorem 4.2** (Hardness of Approximation). *Approximating `OptimalDirectedSearch` up to a  $(\frac{24}{23} - \varepsilon)$ -factor is NP-hard for any  $\varepsilon > 0$ .*

We prove the above hardness result in Section C. From a technical perspective, our approach very closely mirrors that of Chakrabarty and Goel (2010), who show hardness of approximation for related problems (e.g., maximum budgeted allocation) by reducing from Hastad’s 3-bit PCP (Håstad 2001). While their technique applies to our setting, it is not clear that their hardness results apply directly—a challenge specific to our setting is that our optimization problem is continuous rather than discrete. To make this approach work, we must show that our complicated feasibility set introduces a discrete element to the optimal allocation. But as a consequence, we obtain a slightly worse constant  $c$ .

A constant-factor gap between first-best and second-best is also inherent to the `OptimalDirectedSearch` problem. We show via example in Section B.1 that there exists a market such that the first-best and second-best are separated by a factor of 2.

**Proposition 4.3** (Gap Between First- and Second-Best). *For any  $\varepsilon > 0$ , there exists a market such that the platform’s first-best optimal welfare is at least  $2 - \varepsilon$  times the optimal (second-best) welfare for `OptimalDirectedSearch`.*

Before proceeding with the proof of Theorem 4.1, we remark that the first-best optimal welfare and the welfare achieved by random meeting serve as natural benchmarks to compare our solution against. (We remind the reader that, in our general setting, an assortative matching is not well-defined.) Theorem 4.1 ensures that our solution achieves at least 1/4 of the former benchmark. On the other hand, Example 3.4 implies the welfare achieved under our solution can be arbitrarily larger than that achieved by random meeting. In fact, it is straightforward to show that for the

horizontal market of Example 3.4, our solution coincides with the first-best.<sup>18</sup> Consequently, for large  $\delta$ , its welfare gain compared to random meeting would be proportional to  $n$ .

In the remainder of this section, we develop the proof of Theorem 4.1.

#### 4.1. The First-Best Matching Economy

In this section, we study the platform's first-best optimization problem, where the platform has the power to plan the entire matching economy. That is, in addition to designing assortments, we allow the platform to ignore agents' incentives and dictate each type's strategy. The first-best relaxation is obtained by modifying the optimization problem (6) as follows: Relax the best response constraint (6e). Then, replace the per-type threshold variables  $\tau_\theta$  with pairwise threshold variables  $\tau_{mw}$  that take the place of  $\max(\tau_m, \tau_w)$  in all expressions for all  $(m, w) \in \mathcal{M} \times \mathcal{W}$  in (6).

To see why this relaxation corresponds to the platform's first-best problem, note that rather than optimizing over all agent strategies, the platform can restrict its attention to strategies  $\sigma_\theta$  that have a threshold for each type on the opposite side. This is because whenever two agents of types  $m$  and  $w$  match at a certain rate, it is optimal for the matches realized to be the highest-valued ones. In particular, the platform should never hope that agents of types  $m$  and  $w$  reject each other at utility  $u$  but match at some utility  $u' < u$ . Rather, there should exist a threshold  $\tau_{mw}$  such that agents of these types match if and only if  $u \geq \tau_{mw}$ .

While even the relaxed optimization problem looks difficult to work with, our goal in this section is to show that a careful reparametrization reduces the optimization problem to a linear program. The resulting linear program is related to (and in fact generalizes) the linear programming relaxation of the assignment problem; the main difference is that instead of having unit capacities on the flow to each node, we allow for real-valued capacities. Formally, the reduction can be stated as follows:

**Proposition 4.4** (First-Best ‘‘Assignment Problem’’). *Consider the linear program*

$$\max_{\beta_{mw}} 2 \cdot \sum_{m \in \mathcal{M}} \sum_{w \in \mathcal{W}} \rho_{mw} \cdot \beta_{mw} \quad (7)$$

$$\text{such that } \sum_{w \in \mathcal{W}} \beta_{mw} \leq \alpha_m \quad \forall m \in \mathcal{M} \quad (7a)$$

$$\sum_{m \in \mathcal{M}} \beta_{mw} \leq \alpha_w \quad \forall w \in \mathcal{W}, \quad (7b)$$

where  $\rho_{mw} := \max_{\tau \geq 0} \frac{\int_{\tau}^{\infty} u dF_{mw}}{\delta + \int_{\tau}^{\infty} dF_{mw}}$ . The platform's first-best optimization problem is equivalent to (7) in the following sense:

1. The two optimization problems have the same optimal objective value.
2. Any feasible choice of  $\{\beta_{mw}\}_{m \in \mathcal{M}, w \in \mathcal{W}}$  corresponds to a feasible choice of  $\{\lambda_\theta\}_{\theta \in \Theta}$ ,  $\{\eta_\theta\}_{\theta \in \Theta}$ ,  $\{\xi_\theta\}_{\theta \in \Theta}$ , and  $\{\tau_{mw}\}_{m \in \mathcal{M}, w \in \mathcal{W}}$  that achieves the same objective value.

<sup>18</sup> See Section 5 for a numerical solution of a similar horizontal market as well as further numerical analysis.

Proposition 4.4 implies that to solve the first-best optimization problem, it suffices to solve the linear program (7) and convert the resulting solution to an optimal choice of variables for the first-best optimization problem. The proof of the Proposition 4.4 proceeds by first reparameterizing the optimization program (6) using the flow balance constraint: We introduce the variable

$$f_{mw} := \eta_m \lambda_m(w) = \eta_w \lambda_w(m)$$

for the total rate of meetings between each pair  $(m, w) \in \mathcal{M} \times \mathcal{W}$ . We can then encode the flow balance, capacity, and stationary constraints into the single combined constraint

$$\alpha_\theta \geq \sum_{\theta'} \left( f_{\theta\theta'} \left( \delta + \int_{\tau_{\theta\theta'}}^{\infty} dF_{\theta\theta'} \right) \right). \quad (8)$$

That is, there exist corresponding values for  $\lambda_\theta(\theta')$ ,  $\xi_\theta$ , and  $\eta_\theta$  that satisfy the constraints of (6) if and only if the  $f_{\theta\theta'}$  variables satisfy (8). To finish, we introduce the variables

$$\beta_{mw} := f_{mw} \left( \delta + \int_{\tau_{mw}}^{\infty} dF_{mw} \right). \quad (9)$$

We then substitute into (8) and the objective of (6). The latter becomes

$$2 \sum_{m \in \mathcal{M}} \sum_{w \in \mathcal{W}} \left( \beta_{mw} \cdot \frac{\int_{\tau_{mw}}^{\infty} u dF_{mw}}{\delta + \int_{\tau_{mw}}^{\infty} dF_{mw}} \right). \quad (10)$$

While the thresholds  $\tau_{mw}$  are not yet fixed, these substitutions make the optimization effectively unconstrained in  $\tau_{mw}$ . Optimizing over  $\tau_{mw}$  in the objective lets us write the objective in terms of  $\rho_{mw}$  as in (7), since the only dependence on  $\tau_{mw}$  is through the expression

$$\frac{\int_{\tau_{mw}}^{\infty} u dF_{mw}}{\delta + \int_{\tau_{mw}}^{\infty} dF_{mw}} \quad (11)$$

We defer full details of this analysis to Section B.2.


The final step in the above is setting  $\tau_{mw}$  so that expression (11) is maximized. It turns out that we can do so by setting  $\tau_{mw} = \rho_{mw}$ . (That is, the expression achieves its maximum value at a fixed point.) We record this observation here:

**Lemma 4.5** (Fixed Point Structure of Optimal Threshold). *Define*

$$A(\tau) = \frac{\int_{\tau}^{\infty} u dF}{\delta + \int_{\tau}^{\infty} dF},$$

where  $\delta > 0$  and  $F$  is a continuous distribution, and let  $\rho = \max_{\tau \geq 0} A(\tau)$ . Then  $A(\rho) = \rho$ . Moreover,  $A$  is monotonically increasing for  $\tau \leq \rho$  and monotonically decreasing for  $\tau \geq \rho$ .

**Figure 2** This figure depicts the process for converting an optimal first-best solution to a stationary equilibrium, while preserving welfare up to a factor of 4. We start with a first-best solution whose graph of positive weight variables forms a forest (i.e., step (1)). The first arrow represents step (2), where we approximate the forest with stars, each corresponding to a star-shaped submarket. The second arrow represents solving the first-best for each star-shaped submarket (see Lemma 4.9). The third pair of arrows corresponds to the cases we handle in the proof of Proposition 4.8—in particular, whether the new first-best solution induces a stationary equilibrium (Case 1) or not (Case 2).



## 4.2. Designing an Approximately Optimal Equilibrium

In the previous section, we described how the platform can solve for a first-best solution, where agents' incentives are ignored. In this section, we now bring back agents' incentives and show how to convert the first-best solution into a stationary equilibrium. We show that our conversion process retains at least  $\frac{1}{4}$  of the welfare of the first-best solution, thereby giving us an efficient 4-approximation algorithm for `OptimalDirectedSearch`.

At a high-level, our conversion process works as follows (see Figure 2 for a visual depiction):

- Step (1).** Find a solution  $\{\beta_{mw}^*\}_{m \in \mathcal{M}, w \in \mathcal{W}}$  to the linear program (7) such that the graph on  $\Theta$  that has an edge  $(m, w)$  for each  $\beta_{mw}^* > 0$  is a forest.
- Step (2).** Approximate this edge-weighted forest with a union of vertex-disjoint star graphs such that the star graphs contain at least half the total edge weight of the forest.
- Step (3).** Construct a set of assortments for each star-shaped submarket (where one side has only one type of agent) that induces a stationary equilibrium whose social welfare is at least half the first-best social welfare of the submarket.

The first two of the above steps have appeared in other works featuring generalized assignment problems Lenstra et al. (1990), Chakrabarty and Goel (2010), Banerjee et al. (2017), where they were similarly used to simplify the problem structure. Our main insight for this stage of the proof lies in the third step, where we leverage structural properties of the agents' decision problem to show that the optimal stationary equilibrium in each submarket produces at least half the first-best welfare of that submarket.<sup>19</sup>

<sup>19</sup> Note that this last step is distinct from simply computing the optimal stationary equilibrium. Computing the equilibrium is not enough to show that it actually obtains at least half the first-best welfare.

These steps, together with Proposition 4.4, prove our approximation result Theorem 4.1.

**4.2.1. Steps (1) and (2)** We now state the lemmas that give the results needed for steps (1) and (2). The first lemma, having appeared in various forms since Lenstra et al. (1990), shows that there exists an optimal solution for our generalized assignment problem such that the positive variables  $\beta_{mw} > 0$  form a forest when viewed as edges on a graph with vertex set  $\Theta$ . This result is true for the same reason that, in the classical assignment linear program, there exists an optimal solution where all the variables are integral: As long as a cycle exists, one can push “flow” through this cycle to eliminate it while simultaneously reducing the value of some variable to 0.

**Lemma 4.6** (Lenstra et al. (1990), Chakrabarty and Goel (2010), Banerjee et al. (2017)). *For a generalized assignment problem*

$$\begin{aligned} \max \quad & \sum_{m \in \mathcal{M}} \sum_{w \in \mathcal{W}} \rho_{mw} \cdot \beta_{mw} \\ \text{such that} \quad & \sum_{w \in \mathcal{W}} \beta_{mw} \leq \alpha_m \quad \forall m \in \mathcal{M} \\ & \sum_{m \in \mathcal{M}} \beta_{mw} \leq \alpha_w \quad \forall w \in \mathcal{W}, \end{aligned}$$

*with real-valued capacities  $\alpha_\theta$ , there exists an optimal solution  $\{\beta_{mw}^*\}_{m \in \mathcal{M}, w \in \mathcal{W}}$  such that the graph on  $\Theta$  with edge set  $\{(m, w) : \beta_{mw}^* > 0\}$  is a forest.*

The next lemma, which is the main ingredient for step (2), shows that we can take any edge-weighted forest and approximate it up to a factor of 2 in terms of edge weight with a disjoint union of star graphs. The precise formulation we cite is due to Banerjee et al. (2017), though similar ideas have also appeared in the contexts of budgeted allocations Lenstra et al. (1990), Chakrabarty and Goel (2010). This lemma is simple to prove: Root each tree in the forest arbitrarily, and color each edge red or blue based on its distance modulo 2 from its root. The colors partition the forest into two subgraphs each made up of star-shaped graphs; one of these subgraphs must capture at least half the edge weight of the forest.<sup>20</sup>

**Lemma 4.7** (Banerjee et al. (2017)). *Given an edge-weighted forest, there exists a subgraph that is a union of vertex-disjoint star graphs (i.e., trees of radius 1) such that the total edge weight of the subgraph is at least half that of the original tree.*

In context, Lemma 4.6 tells us that we can find a solution  $\{\beta_{mw}^*\}_{m \in \mathcal{M}, w \in \mathcal{W}}$  to the linear program (7) for the platform’s first-best optimization problem such that the positive weight variables form a

<sup>20</sup> We remark that the constant factor on this lemma is tight: Consider the tree with  $2n + 1$  vertices, where the root vertex has  $n$  children and each of these children has another child. Then, any subgraph that is a vertex-disjoint union of star graphs can only obtain edge weight  $n + 1$ .

forest on  $\Theta$ . We weight the edges of this forest by their contribution  $\beta_{mw} \cdot \rho_{mw}$  to social welfare in the first-best optimization problem. Next, Lemma 4.7 tells us that we can 2-approximate this edge-weighted forest with a union of vertex-disjoint star graphs. Each of these star graphs defines a “star-shaped” submarket; note that each such submarket has a side where all agents are of the same type. This completes steps (1) and (2).

**4.2.2. Step (3)** We now show that for any star-shaped submarket (i.e., one where one side of the market consists entirely of agents of the same type), the platform can design a set of assortments that induce a stationary equilibrium with welfare at least half that of the platform’s *first-best* solution. Formally, we prove the following proposition:

**Proposition 4.8** (Star-shaped Markets). *For any star-shaped market, the platform can choose assortments that induce a stationary equilibrium whose welfare is a 2-approximation to the platform’s first-best welfare.*

For the proof of this proposition, a key observation is that the platform’s first-best threshold coincides with an agent’s optimal threshold if their assortment is saturated with agents of a single type. In particular, such agents will not accept more than the platform would want them to in the first-best equilibrium. We obtain from this observation a condition for when the first-best solution induces a stationary equilibrium. We then show that the induced equilibrium achieves at least half the first-best welfare. When the condition fails to hold, it turns out that we can modify the first-best solution appropriately, by only matching agents whose capacity constraint is tight, and still obtain an equilibrium where at least half the first-best welfare is preserved.

For the remainder of this section, we assume without loss of generality that  $\mathcal{M} = \{m\}$ , i.e., our star-shaped market is such that the  $\mathcal{M}$  side of the market consists only of type  $m$  agents.

*Structural Observations for Star-Shaped Markets.* Before giving the proof of Proposition 4.8, we develop some structural observations about the optimal first-best solution and agents’ thresholds. We defer the proofs of all lemmas in this section to Section B.4.

In our first lemma, we characterize the platform’s first-best solution in star-shaped markets:

**Lemma 4.9** (First-best Solution of a Star-Shaped Market). *There exists a first-best solution  $\{\beta_{mw}^*\}_{w \in \mathcal{W}}$  and a subset  $w_1, w_2, \dots, w_n \subseteq \mathcal{W}$  such that  $\rho_{mw_1} \geq \rho_{mw_2} \geq \dots \geq \rho_{mw_n}$ ,  $\beta_{mw_i}^* = \alpha_{w_i}$  for all  $i < n$ , and  $\beta_{mw}^* = 0$  if  $w \notin \{w_1, \dots, w_n\}$ .*

Now that we understand the platform’s first-best optimal solution, we can start to incorporate incentives through the incentive constraint (6e). By Lemma 4.5, each first-best threshold  $\tau_{mw_i}$  in the first-best solution described in Lemma 4.9 can be taken to be  $\rho_{mw_i}$ , since  $\rho_{mw_i}$  maximizes

$$\frac{\int_{\tau_{mw_i}}^{\infty} u dF_{mw_i}}{\delta + \int_{\tau_{mw_i}}^{\infty} dF_{mw_i}}.$$



Then, by the definition  $\beta_{mw_i} = f_{mw_i} \left( \delta + \int_{\tau_{mw_i}}^{\infty} dF_{mw_i} \right)$  in (9), this first-best solution corresponds to the choice of  $f_{mw_i}$  given by

$$f_{mw_i} = f_{mw_i}^* := \frac{\beta_{mw_i}^*}{\delta + \int_{\rho_{mw_i}}^{\infty} dF_{mw_i}}.$$

Our next step is to consider incentives and characterize equilibrium when  $f_{mw_i} = f_{mw_i}^*$ . Fixing the  $f_{mw_i}$  values gives us information about equilibrium thresholds  $\tau_{\theta}$  for all types because the incentive constraint (6e) is equivalent to

$$\tau_{\theta} = \frac{1}{\alpha_{\theta}} \sum_{\theta'} \left( f_{\theta\theta'} \int_{\max(\tau_{\theta}, \tau_{\theta'})}^{\infty} u dF_{\theta\theta'} \right) \quad (12)$$

after scaling both the numerator and the denominator of the right-hand side by  $\eta_{\theta}$ . (Recall that this fixed point definition of equilibrium play comes from Proposition 3.1, which lets us restrict our attention to equilibria where each type has threshold equal to their expected utility.)

As a point of comparison when analyzing agents' equilibrium play, we consider the “idealized thresholds”  $\hat{\tau}_{w_i}$  that arise when type  $w_i$  has absolute market power. These idealized thresholds, which are given by the fixed point equation

$$\hat{\tau}_{w_i} = \frac{1}{\alpha_{w_i}} f_{mw_i}^* \int_{\hat{\tau}_{w_i}}^{\infty} u dF_{mw_i},$$

can also be thought of as the threshold that type  $w_i$  agents would set if  $\tau_m = 0$ . (Note that the left-hand side is monotonically increasing in  $\hat{\tau}_{w_i}$  while the right-hand side is monotonically decreasing, so the fixed point exists and is well-defined.) We care about these idealized thresholds because, in equilibria where  $f_{mw_i} = f_{mw_i}^*$ , the threshold chosen by type  $m$  will be a best response to these idealized thresholds. Namely, type  $m$  will play the threshold  $\hat{\tau}_m$  which satisfies

$$\hat{\tau}_m = \frac{1}{\alpha_m} \sum_{i=1}^n \left( f_{mw_i}^* \int_{\max(\hat{\tau}_m, \hat{\tau}_{w_i})}^{\infty} u dF_{mw_i} \right). \quad (13)$$

To prove this claim, we will proceed through the following series of lemmas. These lemmas and the definition of  $\hat{\tau}_m$  will also be useful for our later analysis.

For the idealized thresholds  $\hat{\tau}_{w_i}$ , it is not difficult to establish using the definition of  $\rho_{mw_i}$  and the monotonicity property in Lemma 4.5 that:

**Lemma 4.10** (Comparison to First-best Thresholds). *If  $i < n$ , then  $\hat{\tau}_{w_i} = \rho_{mw_i}$  (where  $\rho_{mw_i}$  is defined in Proposition 4.4). And at  $i = n$ , it holds that  $\hat{\tau}_{w_n} \leq \rho_{mw_n}$ .*

A similar argument shows that these  $\hat{\tau}_{w_i}$  upper bound the equilibrium thresholds  $\tau_{w_i}$  for all  $i$ :

**Lemma 4.11** (Comparison to Equilibrium Thresholds). *In any equilibrium where  $f_{mw_i} = f_{mw_i}^*$ ,  $\tau_{w_i} \leq \hat{\tau}_{w_i}$  for all  $i$ .*

While the thresholds of type  $w_i$  are always upper bounded by  $\hat{\tau}_{w_i}$ , the quantity  $\max(\tau_m, \tau_{w_i})$  is always at least  $\hat{\tau}_{w_i}$ : Intuitively, if  $\tau_m < \hat{\tau}_{w_i}$ , then type  $w_i$  would set their threshold to be  $\hat{\tau}_{w_i}$ , since that would be optimal for them. This property, together with Lemma 4.10, guarantees that no type  $w_i$  for  $i < n$  will match more than they would have under the platform's first-best solution.

**Lemma 4.12** (Lower Bound on Equilibrium Matching Thresholds). *In any equilibrium where  $f_{mw_i} = f_{mw_i}^*$ ,  $\max(\tau_m, \tau_{w_i}) \geq \hat{\tau}_{w_i}$  for all  $i$ .*

From Lemmas 4.11 and 4.12, we deduce the threshold  $\tau_m$  played in equilibria where  $f_{mw_i} = f_{mw_i}^*$  for all  $i$ : These lemmas together imply that  $\max(\tau_m, \tau_{w_i}) = \max(\tau_m, \hat{\tau}_{w_i})$  in such equilibria. Recall that (by Proposition 3.1)  $\tau_m$  must also satisfy

$$\tau_m = \frac{1}{\alpha_m} \sum_{i=1}^n \left( f_{mw_i}^* \int_{\max(\tau_m, \tau_{w_i})}^{\infty} u dF_{mw_i} \right) = \frac{1}{\alpha_m} \sum_{i=1}^n \left( f_{mw_i}^* \int_{\max(\tau_m, \hat{\tau}_{w_i})}^{\infty} u dF_{mw_i} \right). \quad (14)$$

This is exactly the definition of  $\hat{\tau}_m$  as given in (13), so we may conclude  $\tau_m = \hat{\tau}_m$ .

*Proving the Proposition.* With our structural lemmas, we are now ready to prove Proposition 4.8. We split our analysis into two cases based on the value of  $\hat{\tau}_m$ . We show that if  $\hat{\tau}_m > \rho_{mw_n}$ , then there exists a stationary equilibrium (with incentives) corresponding to the first-best  $f_{mw_i}^*$  values. Moreover, this equilibrium obtains at least half the welfare of the first-best optimal. Otherwise, if  $\hat{\tau}_m \leq \rho_{mw_n}$ , then we consider two further possibilities: If  $\sum_{i=1}^{n-1} \rho_{mw_i} \cdot \alpha_{mw_i}$  is at least half the first-best optimal, then we only let type  $m$  agents meet type  $w_i$  agents for  $i < n$ . Otherwise, we only let type  $m$  agents meet type  $w_n$  agents. In each of these two remaining cases, we will show that we still capture at least half of the first-best optimal welfare.

*Proof of Proposition 4.8.* As mentioned above, we split our analysis into two cases based on the value of  $\hat{\tau}_m$ . For both, we use OPT to refer to the welfare value attained by the platform's first-best optimal solution, i.e.,  $\text{OPT} = 2 \sum_{i=1}^n \rho_{mw_i} \cdot \beta_{mw_i}^*$ . We now discuss the two cases:

**Case 1:  $\hat{\tau}_m > \rho_{mw_n}$ .** If  $\hat{\tau}_m > \rho_{mw_n}$ , we show that the first-best choice of  $f_{mw_i}^*$  induces a stationary equilibrium obtaining at least half the first-best optimal welfare.

To check feasibility, we need to check that the  $\{\xi_\theta\}_{\theta \in \Theta}$  and  $\{\eta_\theta\}_{\theta \in \Theta}$  given by the thresholds in the preceding lemmas satisfy the constraints of (6). By our analysis in Proposition 4.4, it is equivalent to check that the combined constraint (8) holds. By Lemmas 4.10 and 4.12, in equilibrium, we have  $\max(\tau_m, \tau_{w_i}) \geq \hat{\tau}_{w_i} = \rho_{mw_i}$  for all  $i < n$  and  $\tau_m = \hat{\tau}_m$ , so  $\max(\tau_m, \tau_{w_n}) \geq \hat{\tau}_m > \rho_{mw_n}$ . This implies all the thresholds in equilibrium are higher than they were in the platform's first-best solution. In other words, the corresponding matching rates  $\xi_\theta$  are no larger than they were in the first-best solution. It follows that the combined constraint (8) can only have more slack, meaning that a stationary equilibrium is indeed induced by the choice of  $f_{mw_i}^*$  and the thresholds discussed above.

Our remaining task is to check that the realized welfare is sufficient by lower bounding  $\hat{\tau}_m$ . The intuition for the following calculation is that if type  $m$  is rejecting potential matches of value  $u$ , then the expected payoff of type  $m$  must be at least  $u$ . Indeed, we have:

$$\begin{aligned}
\frac{1}{2}\text{OPT} &= \sum_{w \in \mathcal{W}} \rho_{mw} \cdot \beta_{mw}^* \\
&= \sum_{w \in \mathcal{W}} \left( f_{mw}^* \int_{\rho_{mw}}^{\infty} u dF_{mw} \right) \\
&= \sum_{i=1}^n \left( f_{mw_i}^* \int_{\max(\hat{\tau}_m, \rho_{mw_i})}^{\infty} u dF_{mw_i} \right) + \sum_{i=1}^n \left( f_{mw_i}^* \int_{\rho_{mw_i}}^{\max(\hat{\tau}_m, \rho_{mw_i})} u dF_{mw_i} \right) \\
&\leq \sum_{i=1}^n \left( f_{mw_i}^* \int_{\max(\hat{\tau}_m, \hat{\tau}_{w_i})}^{\infty} u dF_{mw_i} \right) + \sum_{i=1}^n \left( f_{mw_i}^* \int_{\rho_{mw_i}}^{\infty} \hat{\tau}_m dF_{mw_i} \right) \\
&\leq \alpha_m \cdot \frac{1}{\alpha_m} \sum_{i=1}^n \left( f_{mw_i}^* \int_{\max(\hat{\tau}_m, \hat{\tau}_{w_i})}^{\infty} u dF_{mw_i} \right) + \hat{\tau}_m \sum_{i=1}^n \beta_{mw_i}^* \\
&\leq 2\alpha_m \hat{\tau}_m.
\end{aligned}$$

The last inequality follows from (13) and the constraint on the  $\beta_{mw}^*$  from (7). Since  $2\alpha_m \hat{\tau}_m$  is the social welfare under the current allocation (where the factor of 2 comes from accounting for the utilities of both sides), we have the desired lower bound on welfare relative to OPT.

**Case 2:**  $\hat{\tau}_m \leq \rho_{mw_n}$ . When  $\hat{\tau}_m \leq \rho_{mw_n}$ , it is possible that the first-best choice of  $f_{mw_i}^*$  may not induce a feasible equilibrium outcome, since type  $w_n$  might accept more than the platform intended in equilibrium, resulting in the combined constraint (8) being violated for type  $m$ . To resolve this, we modify the first-best optimal solution in one of two ways by considering two subcases: In the first subcase, we simply ignore type  $w_n$  and match type  $m$  to types  $w_1, w_2, \dots, w_{n-1}$ . In the second subcase, we only match type  $m$  to type  $w_n$ . At least one of these two subcases will obtain half the first-best optimal welfare.

**Subcase 2(a):**  $\sum_{i=1}^{n-1} \rho_{mw_i} \cdot \beta_{mw_i}^* > \frac{1}{2}\text{OPT}$ . In this subcase, we consider a modified version of the first-best optimal solution from before, where  $\beta_{mw_n}$  is set to 0. We will show that the resulting equilibrium coincides with the first-best equilibrium for the submarket involving only type  $m$  and types  $w_1, \dots, w_{n-1}$ .

First, note that such a modification incurs at most a  $\frac{1}{2}\text{OPT}$  loss in first-best welfare by assumption. Next, recall that type  $m$ 's threshold in equilibrium must satisfy the fixed point equation (13) for the  $f_{mw_i}^*$  values for  $i < n$ . It follows that the resulting  $\tau_m$  is at most  $\hat{\tau}_m$  (e.g., via the argument for Lemma 4.11).<sup>21</sup> Since  $\hat{\tau}_m \leq \rho_{mw_n} \leq \rho_{mw_i} = \hat{\tau}_{w_i}$  for all  $i < n$  by Lemmas 4.9 and 4.10, we must have  $\max(\tau_m, \tau_{w_i}) = \hat{\tau}_{w_i}$  by Lemmas 4.11 and 4.12. This

<sup>21</sup> Here, we clarify that  $\hat{\tau}_m$  is as defined in (13) for the  $f_{mw_i}^*$  values for the *unmodified* first-best market.

shows that the resulting equilibrium is identical to the corresponding first-best solution. (In particular, it is feasible.) Hence the social welfare obtained is  $\sum_{i=1}^{n-1} \rho_{mw_i} \cdot \beta_{mw_i}^* > \frac{1}{2}\text{OPT}$ .

**Subcase 2(b):**  $\rho_{mw_n} \cdot \beta_{mw_n}^* \geq \frac{1}{2}\text{OPT}$ . In this subcase, we consider the first-best solution with  $\beta_{mw} = 0$  for all  $w \neq w_n$ , and set  $\beta_{mw} = \min(\alpha_m, \alpha_{w_n})$ . In this new “single-edged” market, where we only match types  $m$  and  $w_n$ , the first-best social welfare will be at least  $\frac{1}{2}\text{OPT}$  by assumption. In equilibrium, the thresholds must satisfy  $\max(\tau_m, \tau_{w_n}) = \rho_{mw_n}$  by appropriate analogs of Lemmas 4.10 to 4.12. Thus, like in Subcase 2(a), the resulting equilibrium is both feasible and identical to that of the corresponding first-best solution and its social welfare is at least  $\frac{1}{2}\text{OPT}$ .  $\square$

*Wrapping Up.* Theorem 4.1 now follows straightforwardly from our work up to this point. Since the total of the first-best welfares for the each of the “star-shaped” submarkets is at least half the platform’s first-best welfare for the overall market and since Proposition 4.8 shows that at least half of the first-best welfare for each submarket can be realized in stationary equilibrium, we obtain the desired 4-approximation for Theorem 4.1.

## 5. Further Examples via Simulations

To provide further intuition and to give some concrete examples of the model, we present in this section the results of numerical simulations for a set of simple markets. We compare the outcome achieved by our approach to search design against the first-best outcome (defined in Section 4) as well as against the baseline outcome of random meeting.<sup>22</sup> We present two sets of simulations: In the first set of simulations, we interpolate between the two settings—horizontal and vertical—featured in Section 3.3. In the second set of simulations, we highlight the role of search friction (parameterized by  $\delta$ ) on the relative performances of random meeting and designed search.

### 5.1. Setup


For both sets of simulations, the market has four types on each side, i.e.,  $\mathcal{M} = \{m_1, m_2, m_3, m_4\}$  and  $\mathcal{W} = \{w_1, w_2, w_3, w_4\}$ . We assume that  $(\alpha_{m_1}, \alpha_{m_2}, \alpha_{m_3}, \alpha_{m_4}) = (1, 4, 4, 1)$  and  $(\alpha_{w_1}, \alpha_{w_2}, \alpha_{w_3}, \alpha_{w_4}) = (2, 2, 3, 2)$ . Note that these arrival rates are such that there is slight asymmetry between the types on the two sides.<sup>23</sup>

For the first set of experiments, we consider interpolating linearly between a purely horizontal market and a purely vertical market. We assume that each distribution  $\mathcal{D}_{m_i w_j}$  is normal  $\mathcal{N}(\mu_{ij}, \sigma^2)$ ,

<sup>22</sup> While we do not give an algorithm with formal guarantees to compute the outcome of the market under random meeting, the outcome can nonetheless be approximated via fixed point iteration. Specifically, we iteratively compute agents’ thresholds and equilibrium populations, under the assumption that agents meet uniformly randomly, until a fixed point is reached. This iteration-based approximation converges for all markets considered in this section.

<sup>23</sup> This asymmetry leads to more interesting outcomes than if the markets were, say, symmetric. If the markets were symmetric, for our choices of utility values, it would be optimal to fully match type  $m_i$  to type  $w_i$  for all  $i$ .

**Figure 3** On the left-hand side, we plot social welfare for our approximation algorithm and the two benchmarks of the first best and random meeting as  $q$  ranges from 0 to 1. On the right-hand side, we plot the same data, but



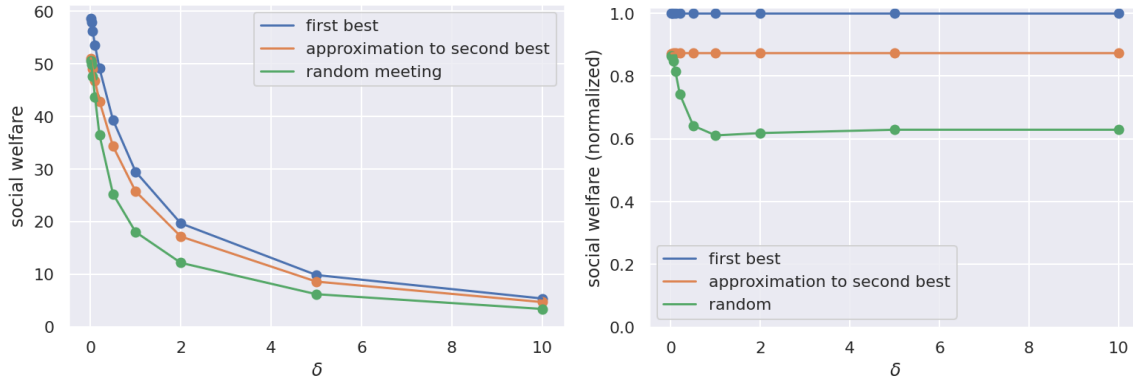
with mean  $\mu_{ij}$  determined linearly interpolating between two extreme horizontal and vertical markets. In the horizontal market, we assume that  $\mu_{ij}^h = 0$  if  $i \neq j$  and  $\mu_{ii}^h = 8$  for all  $i$ . That is, agents only have positive expected utility from a match if they are of “corresponding” types. In the vertical market, we assume that  $\mu_{ij}^v = (5 - i)(5 - j)$  for all  $i, j$ . That is, each type prefers others on the opposite side with smaller indices, with expected utilities being supermodular with respect to this ordering. To interpolate with weight  $q \in [0, 1]$ , we set  $\mu_{ij} = (1 - q)\mu_{ij}^h + q\mu_{ij}^v$ . Finally, we set  $\sigma = 0.1$  and  $\delta = 1$ . The results of these simulations are plotted in Figure 3.

For the second set of experiments, we vary the search friction parameter  $\delta$ . We run this experiment using the market described above, with weighting  $q = 0.5$ . To illustrate the role that  $\delta$  plays in the relative performance of the approaches, we vary it between 0.01 and 10, evaluating the approaches for  $\delta$  belonging to the set  $\{0.01, 0.02, 0.05, 0.1, 0.2, 0.5, 1, 2, 5, 10\}$ . The results of these simulations are plotted in Figure 4.

## 5.2. Discussion

For the first set of simulations (see Figure 3), we find that our approximation algorithm has quite good performance relative to the first best and is better than random meeting—sometimes significantly so—throughout. The performance gap between random meeting and our approximation to the second best is most apparent for the purely horizontal market, corroborating the intuition from Example 3.4 that directed search is particularly important when agents’ preferences have a strong horizontal component. As the market becomes more vertical, the gap between the performances of our approximation algorithm and random meeting narrows. One intuition for why this occurs is that in more vertical markets, while agents match more “assortatively,” they also tend to find more types acceptable and thus random meeting can be more successful relative to the stark profile of the horizontal market. Nonetheless, even in such a setting, there is a clear value to directed search (e.g., via our algorithm).

**Figure 4** On the left-hand side, we plot social welfare for our approximation algorithm and the two benchmarks of the first best and random meeting as  $\delta$  ranges from 0.01 to 10. On the right-hand side, we plot the same data, but normalize the welfare values so that the first-best social welfare is 1 for each market.



For the second set of simulations (see Figure 4), we find that as search friction increases—and thus agents have fewer opportunities to search and match—directed search becomes increasingly necessary to achieve high welfare. Intuitively, such a trend holds because when agents are allowed more opportunities, agents can afford to perform search on their own. On the other hand, when agents only have few opportunities to match, directed search has a larger advantage because the platform can arrange meetings so that agents are able to quickly find desirable matches.

Finally, we note that for the examples considered, our approximation algorithm performs much better than the worst-case 4-factor approximation proved in Theorem 4.1. This is likely because the markets considered for the simulations are relatively simple, with only four types on each side.

## 6. Conclusion

Similar to shopping platforms, matching platforms also rely on recommendation systems to facilitate search by offering personalized assortments. However, the two-sided and decentralized nature of these markets makes the design of their recommendation systems fundamentally different from those used for product recommendation. Congestion and misaligned incentives often necessitate making recommendations that are sub-optimal for certain agents but improve the overall social welfare. In this work, we take a first step toward understanding the intricacies of designing recommendation systems based on imperfect knowledge about preferences while taking agents' strategic behavior into consideration. Somewhat surprisingly, we show that for general symmetric preferences, carefully designed assortments with very limited choice can achieve approximately optimal welfare.

While our framework is general in terms of agent types and the structure of symmetric preferences, it abstracts away from some considerations, which we discuss next. In our model, we assume that once the platform suggests a meeting between two agents, those agents meet, observe the utility of matching with each other, and make a decision to accept or reject the match. However, since

meeting—which corresponds to observing the realized utility via a screening action—can be costly, one might consider a richer model (similar to that of Kanoria and Saban (2021)) in which agents decide whether to (i) accept/reject the match without meeting, (ii) meet at a screening cost and then make an accept/reject decision. Studying search design under such a model for general preference structures is a valuable direction for future work. Incorporating screening cost would involve enriching the set of actions for each agent. This, in turn, would make encoding agents’ strategic behavior in the design problem more complex. Furthermore, the tractability of our framework as well as our design problem crucially relies on having symmetric preferences and a stationary environment. As such, incorporating asymmetric match utilities or non-stationary behavior—such as history-dependent departure rates or strategies—would require new modeling/technical developments and is an interesting direction for future research. It is well-known that even for simple asymmetric preferences and under random meeting, multiple stationary equilibria may arise (Burdett and Coles 1997). Thus, the platform design problem may go beyond designing meeting distributions and involve mechanisms for equilibrium selection.

We conclude by noting that in this work, we focus on the design of matching markets with nontransferable utilities. Such models, also known as marriage models, are motivated by online dating platforms or labor markets with fixed wages. They can also be viewed as a special case of markets with transferable utilities in which the match utility is evenly split between the two parties (Smith 2006). Studying a similar setting with transfers—in which the two agents’ payoffs constitute the Nash bargaining solution—is also well-motivated in the context of online labor markets where wages are negotiable. In such a setting, the platform’s first-best solution remains unchanged. As a future direction, one could explore the possibility of building on our hardness result and approximation scheme to establish analogous results for designing search in matching markets with transferable utilities.

## References

- David J Abraham, Ariel Levavi, David F Manlove, and Gregg O’Malley. The stable roommates problem with globally ranked pairs. *Internet Mathematics*, 5(4):493–515, 2008.
- Heiner Ackermann, Paul W Goldberg, Vahab S Mirrokni, Heiko Röglin, and Berthold Vöcking. Uncoordinated two-sided matching markets. *SIAM Journal on Computing*, 40(1):92–106, 2011.
- Hiroyuki Adachi. A search model of two-sided matching under nontransferable utility. *Journal of Economic Theory*, 113(2):182–198, 2003.
- Philipp Afeche, Rene Caldentey, and Varun Gupta. On the optimal design of a bipartite matching queueing system. *Operations Research*, 2021.
- Mohammad Akbarpour, Shengwu Li, and Shayan Oveis Gharan. Thickness and information in dynamic matching markets. *Journal of Political Economy*, 128(3):783–815, 2020.

- Jerry Anunrojwong, Krishnamurthy Iyer, and Vahideh Manshadi. Information design for congested social services: Optimal need-based persuasion. *Available at SSRN 3849746*, 2021.
- Ali Aouad and Daniela Saban. Online assortment optimization for two-sided matching platforms. *Available at SSRN*, 2020.
- Nick Arnosti and Peng Shi. Design of lotteries and wait-lists for affordable housing allocation. *Management Science*, 66(6):2291–2307, 2020.
- Itai Ashlagi, Patrick Jaillet, and Vahideh H. Manshadi. Kidney exchange in dynamic sparse heterogeneous pools. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce, EC '13*, page 25–26, New York, NY, USA, 2013. Association for Computing Machinery. ISBN 9781450319621. doi: 10.1145/2482540.2482565. URL <https://doi.org/10.1145/2482540.2482565>.
- Itai Ashlagi, Maximilien Burq, Patrick Jaillet, and Vahideh Manshadi. On matching and thickness in heterogeneous dynamic markets. *Operations Research*, 67(4):927–949, 2019a.
- Itai Ashlagi, Anilesh K Krishnaswamy, Rahul Makhijani, Daniela Saban, and Kirankumar Shiragur. Assortment planning for two-sided sequential matching markets. *arXiv preprint arXiv:1907.04485*, 2019b.
- Itai Ashlagi, Mark Braverman, Yash Kanoria, and Peng Shi. Clearing matching markets efficiently: informative signals and match recommendations. *Management Science*, 66(5):2163–2193, 2020.
- Itai Ashlagi, Faidra Monachou, and Afshin Nikzad. Optimal dynamic allocation: Simplicity through information design. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 101–102, 2021.
- Siddhartha Banerjee, Sreenivas Gollapudi, Kostas Kollias, and Kamesh Munagala. Segmenting two-sided markets. In *Proceedings of the 26th International Conference on World Wide Web, WWW 2017, Perth, Australia, April 3-7, 2017*, pages 63–72. ACM, 2017.
- Ken Burdett and Melvyn G. Coles. Marriage and class. *The Quarterly Journal of Economics*, 112(1):141–168, 1997.
- Hector Chade. Matching with noise and the acceptance curse. *Journal of Economic Theory*, 129(1):81–113, 2006.
- Hector Chade, Jan Eeckhout, and Lones Smith. Sorting through search and matching models in economics. *Journal of Economic Literature*, 55(2):493–544, 2017.
- Deeparnab Chakrabarty and Gagan Goel. On the approximability of budgeted allocations and improved lower bounds for submodular welfare maximization and GAP. *SIAM J. Comput.*, 39(6):2189–2211, 2010.
- Yeon-Koo Che and Olivier Tercieux. Optimal queue design. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 312–313, 2021.
- Peter Coles and Ran Shorrer. Optimal truncation in matching markets. *Games and Economic Behavior*, 87: 591–615, 2014.
- Laura Doval and Balász Szentes. On the efficiency of queueing in dynamic matching markets. Technical report, Working paper, Division of the Humanities and Social Sciences, California Institute of Technology, 2019.
- Darrell Duffie, Lei Qiao, and Yeneng Sun. Dynamic directed random matching. *Journal of Economic Theory*, 174:124–183, 2018.



- Ehsan Emamjomeh-Zadeh, Yannai A Gonczarowski, and David Kempe. The complexity of interactively learning a stable matching by trial and error. *arXiv preprint arXiv:2002.07363*, 2020.
- Itai Feigenbaum, Yash Kanoria, Irene Lo, and Jay Sethuraman. Dynamic matching in school choice: Efficient seat reassignment after late cancellations. *Management Science*, 66(11):5341–5361, 2020.
- Yannai A Gonczarowski, Noam Nisan, Rafail Ostrovsky, and Will Rosenbaum. A stable marriage requires communication. *Games and Economic Behavior*, 118:626–647, 2019.
- Hanna Halaburda, Mikolaj Jan Piskorski, and Pinar Yildirim. Competing by restricting choice: The case of matching platforms. *Management Science*, 64(8):3574–3594, 2018.
- Johan Håstad. Some optimal inapproximability results. *J. ACM*, 48(4):798–859, 2001.
- Nicole Immorlica and Mohammad Mahdian. Incentives in large random two-sided markets. *ACM Transactions on Economics and Computation (TEAC)*, 3(3):1–25, 2015.
- Nicole Immorlica, Jacob Leshno, Irene Lo, and Brendan Lucier. Information acquisition in matching markets: The role of price discovery. *Available at SSRN*, 2020.
- Yash Kanoria and Daniela Saban. Facilitating the search for partners on matching platforms. *Management Science*, page to appear, 2021.
- Stephan Lauermaun and Georg Nöldeke. Stable marriages and search frictions. *Journal of Economic Theory*, 151:163–195, 2014.
- Jan Karel Lenstra, David B. Shmoys, and Éva Tardos. Approximation algorithms for scheduling unrelated parallel machines. *Mathematical Programming*, 46:259–271, 1990.
- Lydia T. Liu, Horia Mania, and Michael I. Jordan. Competing bandits in matching markets. In *The 23rd International Conference on Artificial Intelligence and Statistics, AISTATS 2020, 26-28 August 2020, Online [Palermo, Sicily, Italy]*, volume 108 of *Proceedings of Machine Learning Research*, pages 1618–1628. PMLR, 2020.
- Afshin Nikzad. Thickness and competition in ride-sharing markets. *Available at SSRN 3065672*, 2017.
- Adam Ozimek. Report: Freelancing and the economy in 2019. 2019.
- Pew Research Center. The virtues and downsides of online dating. 2020.
- Ignacio Ríos, Daniela Saban, and Fanyin Zheng. Improving match rates in dating markets through assortment optimization. *Available at SSRN*, 2020.
- Alvin E Roth. The economics of matching: Stability and incentives. *Mathematics of operations research*, 7(4): 617–628, 1982.
- Robert Shimer and Lones Smith. Assortative matching and search. *Econometrica*, 68(2):343–369, 2000.
- Robert Shimer and Lones Smith. Matching, search, and heterogeneity. *The BE Journal of Macroeconomics*, 1(1), 2001.
- Lones Smith. The marriage model with search frictions. *Journal of political Economy*, 114(6):1124–1144, 2006.
- Yeneng Sun. The exact law of large numbers via Fubini extension and characterization of insurable risks. *Journal of Economic Theory*, 126(1):31–69, 2006.

## Appendix

### A. Formal Definitions of MDP, Strategies, and Equilibrium

This section develops the mathematical formalization for the agents' decision problem introduced in Section 3.1. We also prove Lemma 3.3 and Proposition 3.1 using the formalization developed here.

#### A.1. Solving for the Agent's Best Response

Recall from Section 3.1 that any fixed set of assortments defines a game for the agents, in which agents decide who to accept and reject in their meetings. Recall also that agents play symmetric time- and history-independent strategies. We thus write the strategy of type  $\theta$  agents as a function  $\sigma_\theta(\theta', u)$  taking values in  $[0, 1]$  which specifies the probability that a type  $\theta$  agent accepts when meeting a type  $\theta'$  agent whom they value at utility  $u$ . Then  $\sigma_\theta(\theta', u) \cdot \sigma_{\theta'}(\theta, u)$  is the probability that two agents of types  $\theta$  and  $\theta'$  who value each other at utility  $u$  mutually accept. For measure-theoretic purposes, we restrict our attention to strategies  $\sigma_\theta$  that are measurable as a function in  $u$  for all  $\theta'$ .

Rather than modeling a full-fledged dynamic game, we model each type  $\theta$  agent as facing a continuous-time MDP defined in terms of their (time-invariant) assortment  $\lambda_\theta$  and opposing types' (time-invariant) strategy profiles  $\sigma_{\theta'}$ . We may do this because, as noted in Section 3.1, agents are only affected by the actions of other types in aggregate; thus any agent's own strategic decision will not affect the (aggregate) action of others. In this MDP, the agent starts in a "waiting" state. From here, they either transition to an "exited" state when they leave unmatched due to a life event or to a "meeting" state when they meet another agent. At this "meeting" state, the agent makes a decision to either "accept" or "reject". Then, if they match, they transition to "exited" with a payoff; otherwise, they transition back to "waiting."

That life events and meetings occur in a memoryless manner actually means we can rid ourselves of the continuous-time aspect of this MDP and get an equivalent discrete-time MDP. Formally, we set up this discrete-time MDP as follows: The agent's initial state is `Waiting`. State `Waiting` transitions to `Exited` with probability  $\delta/(\delta + \sum_{\theta'} \lambda_\theta(\theta'))$  and to `PreMeeting $_{\theta'}$`  with probability  $\lambda_\theta(\theta')/(\delta + \sum_{\theta'} \lambda_\theta(\theta'))$ . State `PreMeeting $_{\theta'}$`  transitions to `Meeting $_{\theta', u}$`  with  $u$  drawn from  $\mathcal{D}_{\theta\theta'}$ . State `Meeting $_{\theta', u}$`  is a decision point where the agent can either `Accept` or `Reject`. If `Accept` is chosen, with probability  $\sigma_{\theta'}(\theta, u)$  they transition to `Exited` and receive payoff  $u$ . In all other cases, they return to state `Waiting`. Finally, `Exited` is a terminal state.

The optimal policy for this MDP admits a simple informal analysis: If  $\bar{u}_\theta$  is the expected payoff of an agent in state `Waiting`, then  $\bar{u}_\theta$  is also the expected continuation payoff if an agent chooses `Reject` in state `Meeting $_{\theta', u}$` . Hence the agent should choose `Accept` if  $u > \bar{u}_\theta$  and `Reject` if  $u < \bar{u}_\theta$ . If  $u = \bar{u}_\theta$ , the agent is indifferent between the two options; however, since the distributions  $\mathcal{D}_{mw}$  are continuous, this occurs with probability 0. We now state this observation more formally:

**Lemma A.1.** *For any best response  $\sigma_\theta$ , let  $\bar{u}_\theta$  be its expected payoff. Then*

$$\sigma_\theta(\theta', u) = \begin{cases} 1 & \text{if } u > \bar{u}_\theta \\ 0 & \text{if } u < \bar{u}_\theta \end{cases}$$

*for all  $\theta'$  such that  $\lambda_\theta(\theta') > 0$  and almost all  $u$  in the support of  $\sigma_{\theta'}(\theta, u) dF_{\theta\theta'}$ .*

Proving such a claim is standard in the analysis of MDPs, so we omit a detailed formal proof.

Lemma A.1 delineates a collection of dominated strategies, namely those that reject potential matches that provide greater utility than the continuation utility (i.e., strategies  $\sigma_\theta$  such that  $\sigma_\theta(\theta', u) < 1$  for some  $\theta'$  and  $u > \bar{u}_\theta$ ). In the subsequent analysis, we rule out such dominated strategies to eliminate “bad” equilibria, e.g., when all agents reject all potential matches (knowing that no one else ever accepts).

Ruling out dominated strategies as above in fact lets us focus on equilibria in *threshold strategies*, where each type has a threshold  $\tau_\theta$  and accepts if and only if  $u \geq \tau_\theta$ . Our next lemma shows that any equilibrium in non-dominated strategies is equivalent, in a certain sense, to an equilibrium in threshold strategies, where each type’s threshold is also their expected utility:

**Lemma A.2.** *Suppose all agents are playing best responses where they accept all potential matches worth at least their expected utility. If type  $\theta$  agents all switch to playing the threshold strategy with threshold  $\tau_\theta = \bar{u}_\theta$ , where  $\bar{u}_\theta$  is the expected utility of type  $\theta$  agents, then each type’s strategy is still a best response. Moreover, the distribution of matches and the utilities at which they are realized remains the same (up to measure 0).*

*Proof.* Since type  $\theta$  agents already accept all potential matches worth at least  $\bar{u}_\theta$ , by switching to the threshold strategy with threshold  $\bar{u}_\theta$ , the only change is that they now reject all potential matches worth less than  $\bar{u}_\theta$ . By Lemma A.1, accepting such potential matches was already a probability 0 event to begin with. Hence the distribution of realized matches remains the same.

To see that all strategies are still best responses, notice that the expected utility of type  $\theta$  agents does not change. Furthermore, type  $\theta$  agents rejecting more matches only restricts the choice sets of agents on the opposite side. Hence their strategies  $\sigma_{\theta'}$  remain best responses also.  $\square$

Applying Lemma A.2 in succession to all types  $\theta \in \Theta$  allows us to convert any equilibrium to an equivalent equilibrium in threshold strategies, such that each agent is thresholding at their expected utility. Thus, without loss of generality, we may restrict our attention to such equilibria.

We can also show a converse of sorts to Lemma A.2, in which we characterize the optimal threshold for type  $\theta$  as a solution to a fixed point equation:

**Lemma A.3.** *There exists a unique fixed point satisfying  $\tau_\theta = \bar{u}_\theta$  (viewing  $\bar{u}_\theta$  as a function of  $\tau_\theta$ ). In particular, if  $\tau_\theta$  satisfies the fixed point equation, then  $\tau_\theta$  is the unique threshold best response for type  $\theta$  agents such that  $\tau_\theta = \bar{u}_\theta$ .*

*Proof.* Let  $\tau_\theta$  be a best response threshold satisfying the fixed point equation. (By Lemma A.2, such a  $\tau_\theta$  exists.) Note that no threshold  $\tau'_\theta > \tau_\theta$  can satisfy the fixed point equation because  $\tau_\theta$  is a best response. So suppose  $\tau'_\theta < \tau_\theta$ . Then, an agent thresholding at  $\tau'_\theta$  would leave the market no later than an agent thresholding at  $\tau_\theta$ . If they leave strictly earlier, then they must have left with payoff at least  $\tau'_\theta$ . And conditioned on leaving at the same time, their expected payoff is  $\tau_\theta$ . The latter occurs with positive probability, so their expected payoff when thresholding at  $\tau'_\theta$  exceeds  $\tau'_\theta$ . It follows that  $\tau'_\theta$  is not a fixed point, making  $\tau_\theta$  the unique fixed point.  $\square$

### A.2. Proof of Lemma 3.3

*Proof of Lemma 3.3.* To compute the agent’s expected payoff at state *Waiting*, we can condition on the event that the agent does not return to *Waiting*. Then, the conditional probability of matching with an agent of type  $\theta'$  is

$$\frac{\lambda_\theta(\theta') \int_{\max(\tau_\theta, \tau_{\theta'})}^{\infty} dF_{\theta\theta'}}{\delta + \sum_{\theta'} \left( \lambda_\theta(\theta') \int_{\max(\tau_\theta, \tau_{\theta'})}^{\infty} dF_{\theta\theta'} \right)}$$

and the expected payoff conditioned on matching with an agent of type  $\theta'$  is

$$\frac{\int_{\max(\tau_\theta, \tau_{\theta'})}^{\infty} u dF_{\theta\theta'}}{\int_{\max(\tau_\theta, \tau_{\theta'})}^{\infty} dF_{\theta\theta'}}.$$

Summing over all  $\theta'$  the product of the above two expression gives us the first equality. An application of Lemma 3.2 along with the definition  $\zeta_\theta = \xi_\theta + \delta$  gives us the second statement.  $\square$

### A.3. Proof of Proposition 3.1: Equilibrium Play for Fixed Assortments

In this section, we complete the proof Proposition 3.1, which states that our model makes a unique prediction of equilibrium play given any fixed set of assortments. We show that there exists a strategy profile where each strategy is a best response and that this strategy profile is unique. Uniqueness implies that stationary equilibria are robustly self-sustaining (e.g., a market in stationary equilibrium cannot be disrupted by agents switching to another strategy profile of best responses) and that the social welfare of an assortment is well-defined (see Section 3.2). Our proof of the existence of a profile of best responses is constructive and can be implemented as an algorithm to compute equilibrium play. Consequently, our model’s unique prediction of equilibrium play can also be efficiently computed by the platform.

To prove that a profile of best responses exists, we show that iterating the best response map converges after finitely many iterations. (In a sense, this iteration behaves like a Gale-Shapley operator.) Our uniqueness result derives from the symmetric valuations assumption that we make (see Section 2.2): Symmetric valuations induce a linear ordering over all possible matches by their valuations. This rules out “cycles” in the preferences and thus the possibility of multiple equilibria.

More formally, let the best response map take as input a strategy profile  $\{\sigma_\theta\}_{\theta \in \Theta}$  and output a new strategy profile  $\{\sigma'_\theta\}_{\theta \in \Theta}$  such that each  $\sigma'_\theta$  is the threshold strategy that thresholds at the expected utility of the agent under the input strategy profile. (This map is a best response by Lemma A.3.) To show existence of equilibrium, we will show that iterating this map on the strategy profile converges.

**Proposition A.4.** *Given assortments  $\{\lambda_\theta\}_{\theta \in \Theta}$ , there is a unique profile  $\{\sigma_\theta\}_{\theta \in \Theta}$  in non-dominated strategies such that each  $\sigma_\theta$  is a best response (up to the equivalence given in Lemma A.2). This strategy profile can be found by iterating the best-response map  $O(|\Theta|)$  times.*

Before we proceed with the proof of Proposition A.4, we note that the claims of Proposition 3.1 follow from combining Lemmas A.2 and A.3 with Proposition A.4.

*Proof.* By applying Lemma A.2, we may restrict our attention to agents playing threshold strategies where each type thresholds at their expected utility  $\bar{u}_\theta$ .

To show existence, we iterate the best response map starting from the strategy profile where each type's threshold is 0. We show that this iteration converges after  $O(|\Theta|)$  steps. In the sequel, when we refer to the thresholds of the  $i$ -th iteration, we mean the thresholds after applying the best response map  $i$  times.

Our first step is to show that each type's threshold has the following pattern: The sequence of thresholds after an odd number of iterations is (weakly) monotonically decreasing and that the sequence of thresholds after an even number of iterations is (weakly) monotonically increasing. We prove this claim by induction. For the base case, note that the claim is trivially true when comparing the zeroth and second iterations, since all thresholds are initialized to 0. Next, suppose the claim is true for iterations  $i - 3$  and  $i - 1$ . Then, note that the thresholds after iterations  $i - 2$  and  $i$  are each given by the expected utilities of agents. Now, if  $i$  is odd, then we know that thresholds after iteration  $i - 1$  are higher than those after iteration  $i - 3$ ; hence, the expected utilities, and therefore the new thresholds, are uniformly lower. The analogous argument works for even  $i$ .

A similar argument shows that for each type, the infimum of that type's thresholds after odd-numbered iterations is at least the supremum of that type's thresholds after even-numbered iterations: The claim clearly holds for the initialization of all thresholds at 0. Next, suppose the claim holds for iteration  $i - 1$ . If  $i$  is odd, then let  $j > i$  be any even number. (Note that the claim follows from the monotonicity property above for  $j < i$ .) The thresholds after the  $j$ -th iteration are best responses to thresholds after the  $(j - 1)$ -th iteration. By the inductive hypothesis, all the thresholds after the  $(j - 1)$ -th iteration are at least those of the  $(i - 1)$ -th iteration. Hence the expected utility of each type in the  $(j - 1)$ -th iteration must be at most that of each type in the  $(i - 1)$ -th iteration. Hence the thresholds after the  $j$ -th iteration are uniformly smaller than the thresholds after the  $i$ -th iteration. The analogous argument works for even  $i$ .

The final step for showing the convergence of this best response iteration is to show that after each odd-numbered iteration, at least one additional type will have their threshold "frozen," meaning that it will not change in any future iterations. Indeed, after any odd-numbered iteration, consider the agent type  $\theta$  with the highest threshold  $\tau_\theta$  who has not yet been shown to be frozen. Then, observe that in all future iterations, no unfrozen agent will have a higher threshold by the two claims above—the first claim handles odd-numbered iterations and the second claim handles even-numbered iterations. It follows that type  $\theta$  would not want to alter its best response in all following iterations: It will not be affected by the strategy of any unfrozen type, because their thresholds will be at most  $\tau_\theta$ ; the strategies of all frozen types will remain the same. Thus, type  $\theta$  will be frozen after this iteration. Since at least one additional type has their threshold frozen after each odd-numbered iteration, iterating the best response map will converge after  $O(|\Theta|)$  iterations.

Next, we show the uniqueness of this equilibrium. Suppose for the sake of contradiction that there are two distinct strategy profiles  $\{\sigma_\theta\}_{\theta \in \Theta}$  and  $\{\sigma'_\theta\}_{\theta \in \Theta}$  such that each strategy is both a best response and a threshold strategy. Let the two sets of thresholds for these two strategy profiles be  $\{\tau_\theta\}_{\theta \in \Theta}$  and  $\{\tau'_\theta\}_{\theta \in \Theta}$ . Furthermore, suppose each threshold is its corresponding type's expected utility. Since the two strategy profiles are distinct, there exists a type  $\theta$  such that  $\tau_\theta \neq \tau'_\theta$  and  $\max(\tau_\theta, \tau'_\theta)$  is maximal. Without loss of generality, we may assume that  $\tau_\theta > \tau'_\theta$ . It follows that if  $\tau_{\theta'} > \tau_\theta$ , then  $\tau_{\theta'} = \tau'_\theta$ , and if  $\tau_\theta \geq \tau_{\theta'}$ , then  $\tau_\theta \geq \tau'_\theta$ .

Consequently,  $\max(\tau_\theta, \tau_{\theta'}) = \max(\tau_\theta, \tau'_{\theta'})$  for all  $\theta'$ . But this gives us the following contradiction: Threshold  $\tau_\theta$  is type  $\theta$ 's expected utility, so Lemma 3.3 tells us

$$\tau'_\theta < \tau_\theta = \frac{\sum_{\theta'} \left( \lambda_\theta(\theta') \int_{\max(\tau_\theta, \tau_{\theta'})}^{\infty} u dF_{\theta\theta'} \right)}{\delta + \sum_{\theta'} \left( \lambda_\theta(\theta') \int_{\max(\tau_\theta, \tau_{\theta'})}^{\infty} dF_{\theta\theta'} \right)} = \frac{\sum_{\theta'} \left( \lambda_\theta(\theta') \int_{\max(\tau_\theta, \tau'_{\theta'})}^{\infty} u dF_{\theta\theta'} \right)}{\delta + \sum_{\theta'} \left( \lambda_\theta(\theta') \int_{\max(\tau_\theta, \tau'_{\theta'})}^{\infty} dF_{\theta\theta'} \right)}.$$

That the right-hand side exceeds the left-hand side means thresholding at  $\tau'_\theta$  is not actually a best response, since we could do better by choosing our threshold to be  $\tau_\theta$ . Therefore, two such distinct strategy profiles cannot exist, proving the proposition.  $\square$

## B. Omitted Proofs from Section 4

### B.1. Proof of Proposition 4.3

*Proof of Proposition 4.3.* We prove this proposition via example. Let  $\varepsilon' = \varepsilon/3$ . Consider a market where  $\mathcal{M} = \{m\}$  and  $\mathcal{W} = \{w_H, w_L\}$ . That is, there is a single type on side  $\mathcal{M}$ , while side  $\mathcal{W}$  is split into agents of high ( $H$ ) and low ( $L$ ) types. We assume that the utility distributions  $\mathcal{D}_{mw_H}$  and  $\mathcal{D}_{mw_L}$  are point masses at  $\frac{1}{\varepsilon'}$  and  $\frac{1}{2}$ , respectively.<sup>24</sup> Suppose further that the arrival rates are such that  $\alpha_m = 2$ ,  $\alpha_{w_H} = \varepsilon'$ , and  $\alpha_{w_L} = 2 - \varepsilon'$ . Finally, we consider this example with  $\delta = \varepsilon'$  (i.e., the market becomes nearly frictionless as  $\varepsilon' \rightarrow 0$ ).

It is not hard to see that the first-best payoff involves matching the entire supply of type  $m$  to the entire supplies of types  $w_H$  and  $w_L$ , with agents accepting all matches. The first-best social welfare is thus

$$\frac{2}{1 + \delta} \left( \alpha_{w_H} \cdot \frac{1}{\varepsilon'} + \alpha_{w_L} \cdot \frac{1}{2} \right) = \frac{4 - \varepsilon'}{1 + \varepsilon'}.$$

Notice that  $\frac{4 - \varepsilon'}{1 + \varepsilon'} \geq 4 - 5\varepsilon'$ .

On the other hand, consider any stationary equilibrium of this market. We consider two cases: whether or not type  $m$  agents match with type  $w_L$  agents. If type  $m$  agents match with type  $w_L$  agents, then the expected utility of type  $m$  agents must be at most  $\frac{1}{2}$ , in which case the total welfare of the market is at most  $2 \cdot \alpha_m \cdot \frac{1}{2} = 2$ . On the other hand, suppose type  $m$  does not match with type  $w_L$  agents. We can upper bound the welfare of this outcome by that of the utility of all type  $w_H$  agents getting matched, i.e.,  $2 \cdot \alpha_{w_H} \cdot \frac{1}{\varepsilon'} = 2$ . It follows that the ratio between the first-best welfare and the best possible outcome for `OptimalDirectedSearch` is at least

$$\frac{\frac{4 - \varepsilon'}{1 + \varepsilon'}}{2} \geq 2 - \frac{5}{2}\varepsilon' > 2 - \varepsilon'.$$

$\square$

### B.2. Proof of Proposition 4.4

*Proof of Proposition 4.4.* We first reparameterize the flow balance constraint by introducing new variables

$$f_{mw} := \eta_m \lambda_m(w) = \eta_w \lambda_w(m)$$

<sup>24</sup> Technically, these distributions do not satisfy our continuity assumption. We present our example this way for simplicity's sake—the same example can be made to work with continuous distributions by adding a small amount of Gaussian noise to each utility value.

to replace  $\lambda_m(w)$  and  $\lambda_w(m)$ . (For symmetry in notation, we will use  $f_{mw}$  and  $f_{wm}$  interchangeably.) With these new variables, we may rewrite the objective symmetrically as

$$2 \sum_{m \in \mathcal{M}} \sum_{w \in \mathcal{W}} \left( f_{mw} \int_{\tau_{mw}}^{\infty} u dF_{mw} \right).$$

Furthermore, note that  $\eta_\theta > 0$  by (6a). Hence we may scale the capacity constraint (6c) by  $\eta_\theta$  and replace all occurrences of  $\eta_\theta \lambda_\theta(\theta')$  with  $f_{\theta\theta'}$  to obtain the equivalent constraint

$$\eta_\theta \geq \sum_{\theta'} f_{\theta\theta'} \quad (15)$$

for all  $\theta \in \Theta$ . We similarly replace the matching rate equation (6d) with

$$\xi_\theta \eta_\theta = \sum_{\theta'} \left( f_{\theta\theta'} \int_{\tau_{\theta\theta'}}^{\infty} dF_{\theta\theta'} \right). \quad (16)$$

Our next step is to absolve ourselves of the variables  $\xi_\theta$  and  $\eta_\theta$ . For the former, we can simply merge constraints (6a) and (16), as both are equalities involving  $\xi_\theta \eta_\theta$ , to get

$$\alpha_\theta = \delta \eta_\theta + \sum_{\theta'} \left( f_{\theta\theta'} \int_{\tau_{\theta\theta'}}^{\infty} dF_{\theta\theta'} \right). \quad (17)$$

Now, from (17), we can get a definition of  $\eta_\theta$ , which we can substitute into the only remaining constraint (15) on  $\eta_\theta$  to get a combined flow balance, capacity, and stationarity constraint

$$\alpha_\theta \geq \sum_{\theta'} \left( f_{\theta\theta'} \left( \delta + \int_{\tau_{\theta\theta'}}^{\infty} dF_{\theta\theta'} \right) \right) \quad (18)$$

for each  $\theta \in \Theta$ .

To obtain the linear program constraints, we make one final substitution of

$$\beta_{mw} := f_{mw} \left( \delta + \int_{\tau_{mw}}^{\infty} dF_{mw} \right), \quad (19)$$

into (18), which gives us constraints (7a) and (7b) of the linear program. We also substitute  $\beta_{mw}$  into the objective, which yields the equivalent objective

$$2 \sum_{m \in \mathcal{M}} \sum_{w \in \mathcal{W}} \left( \beta_{mw} \cdot \frac{\int_{\tau_{mw}}^{\infty} u dF_{mw}}{\delta + \int_{\tau_{mw}}^{\infty} dF_{mw}} \right). \quad (20)$$

The only remaining variables that we haven't taken into account are the thresholds  $\tau_{mw}$ . However, notice that the optimization problem is effectively unconstrained in  $\tau_{mw}$ —given any choice of  $\beta_{mw}$  and  $\tau_{mw}$ , we may choose a  $f_{mw}$  so that the definition (19) is satisfied. Thus, we may simply set  $\tau_{mw}$  so that its contribution

$$\frac{\int_{\tau_{mw}}^{\infty} u dF_{mw}}{\delta + \int_{\tau_{mw}}^{\infty} dF_{mw}} \quad (21)$$

to the objective is maximized. As this is exactly the definition of  $\rho_{mw}$ , we conclude that the linear program (7) has the same objective value as the platform's first-best optimization problem.

It remains to prove the second claim in the equivalence between the linear program (7) and the platform's first-best optimization problem. For this, we note that given any feasible choice of  $\{\beta_{mw}\}_{m \in \mathcal{M}, w \in \mathcal{W}}$ , we can reverse all of the substitutions made and obtain corresponding values for  $\{\lambda_\theta\}_{\theta \in \Theta}$ ,  $\{\eta_\theta\}_{\theta \in \Theta}$ ,  $\{\xi_\theta\}_{\theta \in \Theta}$ , where the variables  $\{\tau_{mw}\}_{m \in \mathcal{M}, w \in \mathcal{W}}$  are set as to maximize (21).  $\square$

### B.3. Proof of Lemma 4.5

*Proof of Lemma 4.5.* The logarithmic derivative of  $A(\tau)$  is

$$\frac{d}{d\tau} \log(A(\tau)) = \frac{F'(\tau)}{\delta + \int_{\tau}^{\infty} dF} - \frac{\tau F'(\tau)}{\int_{\tau}^{\infty} u dF} = \left( \int_{\tau}^{\infty} (u - \tau) dF - \tau \delta \right) \frac{F'(\tau)}{(\delta + \int_{\tau}^{\infty} dF) \left( \int_{\tau}^{\infty} u dF \right)}.$$

The sign of this derivative is given by the first term on the right-hand side, since the fraction on the right-hand side is always non-negative. This first term is monotonically decreasing:

$$\frac{d}{d\tau} \left( \int_{\tau}^{\infty} (u - \tau) dF - \tau \delta \right) = - \left( \delta + \int_{\tau}^{\infty} dF \right) < 0.$$

Therefore, the expression is maximized when this term vanishes, i.e., when

$$\tau = \frac{\int_{\tau}^{\infty} u dF}{\delta + \int_{\tau}^{\infty} dF} = \rho.$$

Moreover, since the derivative is non-negative when  $\tau \leq \rho$  and is non-positive when  $\tau \geq \rho$ , we have the desired monotonicities when  $\tau \leq \rho$  and  $\tau \geq \rho$  as well.  $\square$

We also state and prove a generalization of Lemma 4.5 to non-zero thresholds and multiple types on the opposite side. That is, the fixed point and monotonicity properties hold generally:

**Lemma B.1.** *Define*

$$B(\tau) = \frac{\sum_{i=1}^n \left( \lambda_i \int_{\max(\tau, \tau_i)}^{\infty} u dF_i \right)}{\delta + \sum_{i=1}^n \left( \lambda_i \int_{\max(\tau, \tau_i)}^{\infty} dF_i \right)},$$

where  $\delta > 0$ ,  $\lambda_i > 0$  for all  $i$ ,  $\tau_i \geq 0$  for all  $i$ , and  $F_i$  is a continuous distribution for each  $i$ . Let  $\rho$  satisfies the fixed point equation  $\rho = \max_{\tau \geq 0} B(\tau)$ . Then  $B(\rho) = \rho$ . Moreover,  $B$  is monotonically increasing for  $\tau \leq \rho$  and monotonically decreasing for  $\tau \geq \rho$ .

*Proof.* Note that we can write  $B(\tau)$  as

$$B(\tau) = \frac{\int_{\tau}^{\infty} u dG}{\delta + \int_{\tau}^{\infty} dG},$$

where  $dG = \sum_{i=1}^n (\lambda_i \cdot \mathbf{1}_{\geq \tau_i} \cdot dF_i)$  is a new measure defined on  $\mathbb{R}$  in terms of  $\lambda_i$ ,  $\tau_i$ , and  $F_i$ . Here,  $\mathbf{1}_{\geq \tau_i}$  denotes the indicator function that is 1 on inputs at least  $\tau_i$  and 0 otherwise. After appropriately normalizing the numerator and denominator (by  $\int_{\mathbb{R}} dG$ ), the claim follows from Lemma 4.5.  $\square$

### B.4. Proofs of Lemmas from Section 4.2.2

*Proof of Lemma 4.9.* Recall the platform's first-best optimization problem (7). In a star-shaped market, this linear program reduces to a fractional knapsack problem where the platform has a knapsack of size  $\alpha_m$  and wishes to fill it with items for each  $w \in \mathcal{W}$  of weight  $\alpha_w$  and value  $\rho_{mw}$ . Thus, there exists a first-best solution  $\{\beta_{mw}^*\}_{w \in \mathcal{W}}$  that matches type  $m$  agents exclusively with some subset  $\{w_1, \dots, w_n\} \subseteq \mathcal{W}$  such that  $\rho_{mw_1} \geq \rho_{mw_2} \geq \dots \geq \rho_{mw_n}$  and  $\beta_{mw_i}^* = \alpha_w$  for all  $i < n$ .  $\square$

*Proof of Lemma 4.10.* If  $i < n$ , then  $\beta_{mw_i}^* = \alpha_{w_i}$ . Therefore,

$$\hat{\tau}_{w_i} = \frac{1}{\alpha_{w_i}} f_{mw_i}^* \int_{\hat{\tau}_{w_i}}^{\infty} u dF_{mw_i} = \frac{\int_{\hat{\tau}_{w_i}}^{\infty} u dF_{mw_i}}{\delta + \int_{\rho_{mw_i}}^{\infty} dF_{mw_i}}.$$



And by Lemma 4.5,  $\hat{\tau}_{w_i} = \rho_{mw_i}$  solves the equation. For  $i = n$ , since  $\beta_{mw_n}^* \leq \alpha_{w_n}$ , we have

$$\hat{\tau}_{w_n} = \frac{1}{\alpha_{w_n}} f_{mw_n}^* \int_{\hat{\tau}_{w_n}}^{\infty} u dF_{mw_n} \leq \frac{\int_{\hat{\tau}_{w_n}}^{\infty} u dF_{mw_n}}{\delta + \int_{\rho_{mw_n}} dF_{mw_n}}.$$

If  $\hat{\tau}_{w_n} > \rho_{mw_n}$ , then by Lemma 4.5, we would have that the right-hand side is at most  $\rho_{mw_n}$ , since the expression  $\int_{\tau}^{\infty} u dF_{mw_n}$  is monotonically decreasing in  $\tau$ . This contradicts the assumption that  $\hat{\tau}_{w_n} > \rho_{mw_n}$ . Hence it must be the case that  $\hat{\tau}_{w_n} \leq \rho_{mw_n}$ .  $\square$

*Proof of Lemma 4.11.* This follows from the fact that for any threshold  $\tau_m$ , we have that

$$\tau_{w_i} = \frac{1}{\alpha_{\theta}} f_{mw_i}^* \int_{\max(\tau_m, \tau_{w_i})}^{\infty} u dF_{mw_i} \leq \frac{1}{\alpha_{w_i}} f_{mw_i}^* \int_{\tau_{w_i}}^{\infty} u dF_{mw_i}.$$

The argument used in Lemma 4.10 when  $i = n$  then lets us conclude that  $\tau_{w_i} \leq \hat{\tau}_{w_i}$  for all  $i$ .  $\square$

*Proof of Lemma 4.12.* The lemma is clearly true if  $\tau_m \geq \hat{\tau}_{w_i}$ . Now, suppose  $\tau_m < \hat{\tau}_{w_i}$ . Then  $\hat{\tau}_{w_i}$  would satisfy the fixed point equation (12), meaning  $\hat{\tau}_{w_i}$  would be the equilibrium threshold of type  $w_i$ .  $\square$

### C. Hardness of Approximation

In this section, we show that the platform's computational problem of finding a  $c$ -approximate welfare-maximizing stationary equilibrium is NP-hard for some constant  $c > 1$ . This shows that, conditioned on  $P \neq NP$ , one cannot hope to obtain better than a constant factor approximation to the optimal in polynomial time.

From a technical perspective, our approach very closely mirrors that of Chakrabarty and Goel (2010), who show hardness of approximation for related problems (e.g., maximum budgeted allocation). While their technique applies to our setting, it is not clear that their hardness results apply directly—a challenge specific to our setting is that our optimization problem is continuous rather than discrete. To make this approach work, we must show that our complicated feasibility set introduces a discrete element to the optimal allocation. But as a consequence, we obtain a slightly worse constant  $c$ .

At a high level, we show hardness of approximation by reducing the problem of approximating MAX3LIN2 to finding a sufficiently good approximation for the platform's welfare maximization problem. (Recall that MAX3LIN2 is the optimization problem where, given  $m$  linear equations in  $n$  variables over  $\mathbb{F}_2$  such that each equation involves exactly three variables, the objective is to maximize the number of equations simultaneously satisfied by an assignment to the variables.) That MAX3LIN2 is hard to approximate is due to Håstad (2001) and is a direct corollary of his celebrated 3-bit PCP. We state this hardness result in terms of a promise problem version of MAX3LIN2:

**Theorem C.1** (Håstad (2001)). *It is NP-hard to distinguish between MAX3LIN2 instances where all but  $\varepsilon m$  equations are simultaneously satisfiable and where at most  $(\frac{1}{2} + \varepsilon)m$  equations are simultaneously satisfiable for any  $\varepsilon > 0$ .*

From this classical starting point, we derive hardness of approximation for the platform's optimization problem:

**Theorem C.2.** *Computing a  $(\frac{24}{23} - \varepsilon)$ -approximate solution to the platform's welfare maximization problem is NP-hard for any  $\varepsilon > 0$ .*

*Proof.* As stated above, we reduce instances of MAX3LIN2 to instances of the platform's welfare maximization problem. Specifically, we reduce to instances where the discount rate  $\delta$  is small and the distributions  $\mathcal{D}_{mw}$  are point masses.<sup>25</sup>

Consider a MAX3LIN2 instance in  $n$  variables  $x_1, x_2, \dots, x_n$  and  $m$  equations  $x_{i_\ell} \oplus x_{j_\ell} \oplus x_{k_\ell} = b_\ell$  for  $\ell \in [m]$ . We construct an instance of our matching market by defining types, arrival rates, and payoffs based on these variables and equations:

**Types.** Define *variable types*  $x_{i,0}, x_{i,1} \in \mathcal{M}$  and *switch type*  $s_i \in \mathcal{W}$  for each variable  $x_i$ . Define *equation types*  $e_{\ell,(0,0,0)}, e_{\ell,(0,1,1)}, e_{\ell,(1,0,1)}, e_{\ell,(1,1,0)} \in \mathcal{W}$  for each equation  $x_{i_\ell} \oplus x_{j_\ell} \oplus x_{k_\ell} = b_\ell$ , one type for each satisfying assignment to equation  $\ell$ . That is, type  $e_{\ell,(b_i,b_j,b_k)}$  corresponds to the satisfying assignment  $(x_{i_\ell}, x_{j_\ell}, x_{k_\ell}) = (b_i \oplus b_\ell, b_j \oplus b_\ell, b_k \oplus b_\ell)$ .

**Arrival rates.** Let the variable and switch types for  $x_i$  each arrive at rate  $\alpha_i = 4m_i$ , where  $m_i$  is the number of equations that involve  $x_i$ . Let each equation type arrive at rate 3.

**Payoffs.** Let matching type  $x_{i,b}$  and type  $s_i$  yield payoff  $2(1 + 2\delta)$  for all  $b$ . Let matching type  $x_{i,b}$  with any equation type whose corresponding satisfying assignment has the variable  $x_i$  with parity  $b$  yield payoff 1. All other matches provide no utility.

With this setup, no agent on side  $\mathcal{W}$  will ever reject positive utility matches in equilibrium, since all positive utility matches yield the same payoff for them. Moreover, no switch type agent will ever be rejected in equilibrium, since they yield maximal payoff for variable type agents.<sup>26</sup> And if in equilibrium an agent of type  $x_{i,b}$  has expected continuation utility greater than 1, then they must only accept agents of type  $s_i$ . These observations set the stage for our key lemma:

**Lemma C.3.** *In any welfare-maximizing equilibrium, the expected payoff of all switch type agents will be  $2(1 + 2\delta)/(1 + \delta)$ . Furthermore, such an equilibrium can be achieved by only showing type  $s_i$  and type  $x_{i,b}$  agents to each other for some  $b \in \{0, 1\}$  for each  $i$ .*

*Proof.* First, we show that some variable type  $x_{i,b}$  must have expected payoff greater than 1 for each  $i \in [n]$ . Suppose otherwise, that there exists a welfare-maximizing stationary equilibrium such that the expected payoffs of both variable types  $x_{i,b}$  are at most 1 for some  $i$ . Then we could stop matching both types entirely, by setting their assortments to 0. This comes at a loss in welfare of at most  $2 \cdot 4m_i$ . Then, we could match type  $x_{i,0}$  entirely to  $s_i$ , and obtain an increase in welfare of  $4m_i \cdot 2(1 + 2\delta)/(1 + \delta) > 2 \cdot 4m_i$ . Finally, note that these new assortments also produce a stationary equilibrium: Unmatching the variable types  $x_{i,b}$  only adds slack to the assortment feasibility constraints. Unmatching also does not affect any equation type's threshold,

<sup>25</sup> Although point masses are not continuous distributions, which we required in Section 2.2, we can make point masses continuous by adding a small amount of (bounded) noise. Doing so does not change any of the conclusions. To simplify our exposition, we focus on the case of point masses only.

<sup>26</sup> When  $\delta > 0$ , these observations still hold for distributions that are not point masses, so long as the variation in payoff is sufficiently small.

since they always accept in the first place. This new equilibrium has higher welfare, which contradicts our premise that the original equilibrium is welfare-maximizing.

The preceding argument shows that in any welfare-maximizing equilibrium, there is a variable type  $x_{i,b}$  whose expected utility is greater than 1 for each  $i$ . Notice that this type must exclusively match with type  $s_i$  agents, since equation type agents only yield payoff 1 and thus cannot be accepted in equilibrium. To prove the second part of this lemma, we show that we can obtain another welfare-maximizing equilibrium by matching the entire supply of type  $s_i$  agents to this variable type  $x_{i,b}$ .

Suppose we have a welfare-maximizing equilibrium, and we modify it by matching the entire supply of type  $s_i$  agents to a variable type  $x_{i,b}$  whose expected utility exceeds 1. To restore a stationary equilibrium, note that making this modification may violate two sets of constraints: First, it might violate type  $x_{i,b}$ 's feasibility constraint; but this is not a problem, since we can simply stop showing  $x_{i,b}$  type agents any types but type  $s_i$ , since these agents were only accepting type  $s_i$  agents in the first place. Second, it might cause type  $x_{i,b\oplus 1}$  to want to accept more equation type agents because they stop matching with switch type agents entirely. To resolve this, we can simply reduce the supply of equation type agents to type  $x_{i,b\oplus 1}$  so that by accepting all such agents,  $x_{i,b\oplus 1}$  accepts the same quantity of agents as in the original equilibrium. Finally, to see that this new equilibrium has the same welfare as the original equilibrium, note that the flow rates of matches for both switch and equation types are at least as high as before.  $\square$

This lemma tells us that there exists a welfare-maximizing equilibrium where, for every  $i$ , one of variable types  $x_{i,0}$  and  $x_{i,1}$  is fully matched to the switch type  $s_i$ . Moreover, it suffices to optimize over such equilibria. Notice that each such equilibrium also corresponds to an assignment  $(x_1, \dots, x_n) \in \{0, 1\}^n$ , where  $x_i = b$  if type  $x_{i,b\oplus 1}$  is fully matched to the switch type  $s_i$ .

We are now ready to state our analogs of Lemmas 4.6 and 4.7 of Chakrabarty and Goel (2010):

**Lemma C.4.** *Given a MAX3LIN2 instance with  $m' \geq (1 - \varepsilon)m$  simultaneously satisfiable equations, the platform's problem it maps to has  $\text{OPT} \geq (36 + 48\delta - 12\varepsilon)m/(1 + \delta)$ .*

*Proof.* Given an assignment  $(x_1, x_2, \dots, x_n) \in \{0, 1\}^n$  that satisfies  $m'$  equations, we can match the switch types to the variable types as above. We can then match the remaining variable types to the equation types so that each variable type matches with the equation types of each satisfied equation it is part of at a flow rate of  $4/(1 + \delta)$ . We thus have payoff  $2(1 + 2\delta)/(1 + \delta) \cdot \sum_i 4m_i$  from switch types and payoff  $12/(1 + \delta)$  per satisfied equation. Summing these quantities gives us the lemma.  $\square$

**Lemma C.5.** *Given a MAX3LIN2 instance with  $m' \leq (\frac{1}{2} + \varepsilon)m$  simultaneously satisfiable equations, the platform's problem it maps to has  $\text{OPT} \leq (34.5 + 48\delta + 3\varepsilon)m/(1 + \delta)$ .*

*Proof.* Take any assignment  $(x_1, x_2, \dots, x_n) \in \{0, 1\}^n$  and match switch types to variable types as described above. Now, for any equation that is not satisfied, notice that there must be some equation type that cannot be matched at all to any of the remaining variable types. So the payoff from the equation types for unsatisfied equations is at most  $9/(1 + \delta)$ . Upper bounding the payoffs of the equation types of the satisfied equations by  $12/(1 + \delta)$  per satisfied equation, we get  $12m \cdot 2(1 + 2\delta)/(1 + \delta)$  from switch types and at most payoff  $((\frac{1}{2} - \varepsilon)m \cdot 9 + (\frac{1}{2} + \varepsilon)m \cdot 12)/(1 + \delta)$  from equation types. Summing these quantities proves the lemma.  $\square$

To finish, observe that Lemmas C.4 and C.5 show that if we can approximate the platform's welfare-maximization problem with a factor of

$$\frac{36 + 48\delta - 12\varepsilon}{34.5 + 48\delta + 3\varepsilon},$$

then we would be able to distinguish the two classes of MAX3LIN2 inputs in Theorem C.1. Thus, by taking sufficiently small  $\delta$  and  $\varepsilon$ , we see that it is NP-hard to approximate the platform's welfare-maximization problem within any constant factor less than  $\frac{36}{34.5} = \frac{24}{23}$ .  $\square$