Matching Supply and Demand with Mismatch-Sensitive Players

Mingliu Chen
The Department of Industrial Engineering and Operations Research, Columbia University. mc5006@columbia.edu

Peng Sun
The Fuqua School of Business, Duke University. psun@duke.edu

Zhixi Wan
Faculty of Business and Economics, University of Hong Kong. zhixiwan@hku.hk

We study matching over time with short- and long-lived players who are very sensitive to mismatch. To characterize the mismatch, we model players' preferences as uniformly distributed on a circle, so the mismatch between two players is characterized by the one-dimensional circular angle between them. This framework allows us to capture matching processes in applications ranging from ride sharing to job hunting. Our analytical framework relies on threshold matching policies. If the matching process is controlled by a central planner (e.g. an online matching platform), the matching threshold reflects the trade-off between matching rate and matching quality. We further compare the centralized system with decentralized systems, where players choose their matching partners. We find that matching controlled by either side of the market may achieve optimal or near optimal social welfare, but have great impact on welfare allocation. In particular, letting long-lived players choose their matching partner leaves short-lived players with zero surplus. Moreover, we extend our model with player heterogeneity. Letting long-lived players choose their matching partner leads to better social welfare when the market of short-lived players is thick and the level of heterogenity is significant. Otherwise, letting short-lived players choose matching partners is better.

Key words: Circular City, Individual Preferences, Market Design, Matching, Queueing Theory.
History: Current version: October 2, 2020.

---

1. Introduction

Recent years witness the emergence and rapid growth of a variety of online platforms that match two-sided market players in a timely fashion. Take Didi Hitch, one of the world's largest commuter carpooling platforms, as our motivating example. During morning (or evening) rush hours in hundreds of cities in China, about one million of riders and drivers post individual destinations on the platform and look for their matches. Successful matchings mostly take place in five to twenty minutes, because both sides of the users participate under time constraints. In particular, riders usually have access to a variety of alternative transportation options (e.g., taxi, publication transportation), tending to actively look for outside options and leave the carpooling platform soon unless finding a good match. Driver users, mostly commuters who look for matches to offset gasoline costs, tend to be more patient but also need to depart eventually after some time. Note that this carpooling platform is unlike ride-hailing platforms such as Uber and Lyft, which run centralized algorithms to assign a stream of riders' orders to a pool of online drivers who typically have no stringent destinations.

Obviously, players prefer to be matched with others who have similar destinations. In some cases, especially morning rush hours, they exhibit high sensitivity to matching quality and would decline a matching proposal if that match would cause unbearable detours, which means extra travel time and uncertainties. The platform cares about both matching rate and quality. The former determines the transaction volume and the platform's own short-term revenue, and the latter affects users' experience and their long-term retention rates.

To facilitate matching, the platform has mainly two instruments. One is pricing and the other is its matching mechanism. Once a rider submits her request information, which primarily includes her origin, destination, and departure time preference, the platform quotes a price based on a pre-specified formula that considers the rider's route and time information. If the rider accepts the price, the platform tries to match her with a driver who has an active offer. A driver offer becomes active when he posts its intention to carpool together with his origin, destination, and time preferences; the offer becomes inactive once the driver is matched with a rider, or he withdraws the offer and leaves the platform (e.g., departs without finding a good match). The platform's internal research finds that, also quite intuitively, higher price improves drivers' tolerance of matching quality (i.e., the degree of detours); thus, once the price is accepted by the rider, it potentially affects the driver's willingness to accept the match.

The matching mechanism requires both sides' mutual acceptance. However, the platform can control the process to effectively determine who selects whom. The status quo (at the time when our project started) gave the driver side unlimited chance to forgo matchable orders, and this seemed to have caused drivers to behave picky, namely, to turn down the platform's proposal with matchable riders in the hope for an even better match. Alternatively, the platform was considering curbing drivers' picky behavior by requiring them to specify a maximum acceptable detour at the time of submitting their offers and commit to accept any match within the detour. With drivers' commitment, the platform would invite each rider to make a selection by showing a list of drivers that the rider's destination is within their maximum acceptable detours.

The platform was interested in understanding: first, how the choice of the matching mechanism would affect matching quantity and quality, as well as both sides' users' satisfaction; second, how to jointly optimize pricing and matching mechanism. We attempt to address these two questions in this paper.

We model a two-sided matching market where both sides of players arrive in Poisson processes. Building upon the observation from DiDi Hitch, one side of players (riders) are short-lived, who leave the platform immediately if not being matched upon arrival. The other side of players (drivers) are long-lived, who stay on the platform for exponentially distributed periods of time without a match. A match occurs if and only if both sides accept it. The platform determines a *price* that a short-lived player (i.e., rider) needs to pay to a long-lived player (i.e., driver) in each match.

A salient feature of our model is that we provide a novel method to construct a measure of mismatch between players. In particular, all incoming players' preferences are assumed to be uniformly distributed on the boundary of a circle. The degree of mismatch between two players is characterized by the one dimensional circular angle between them, which we refer to as the *mismatch angle*. We use this one-dimensional scalar to describe the compatibility between any two players, and we focus on *threshold policies* that match two players if their mismatch angle is small enough. This is natural in ride sharing and other scenarios with spatial features. For example, it captures the difference between a rider's destination and a driver's destination, which causes detours in the context of DiDi Hitch. In the carpooling example, it is mainly drivers that bear mismatch angles, given that riders are typically dropped off at their requested destinations.

Note that using circular model to characterize mismatch between players is not restrictive to the carpooling context. For example, on a gig job hunting platform, service seekers want to find suitable candidates as soon as possible. A candidate may be more patient but will not settle with a job requiring very different skills than what she has already possessed. A mismatch angle in our model captures the difference between the skill a candidate possesses and the talent a seeker is looking for. The smaller the angle, the more compatible two players.

Our main model and its analysis (Sections 4 and 5, respectively) focus on the scenario where riders' outside options are homogeneous; in the carpooling context, this happens when only one outside option (e.g., taxi) is available or dominates on the market. In Section 6, we examine the extension where multiple outside options exist, so riders have heterogeneous preferences. To parsimoniously capture the platform's interests in both matching quantity and quality, we use social welfare, defined by the total surplus accrued to riders and drivers, as the platform's primary objective when it designs pricing, as well as the centralized matching mechanism.

In both scenarios (Sections 5 and 6), we first derive the centralized solution benchmark in which the platform optimizes both the pricing and the matching mechanism, and then we study two decentralized matching mechanisms. One of them, referred to as the long-lived-select mechanism, models the status-quo in which the platform determines pricing but allows long-lived players (i.e., drivers) to select short-lived ones (i.e., riders). The other, referred to as the short-lived-select mechanism, models the alternative mechanism, under which the platform determines pricing but effectively allows short-lived players (i.e., riders) select the best match whenever it is matchable.

Our solution of the optimal centralized mechanism sheds light to the trade off between matching quantity and quality. In particular, comparing the optimal mechanism with the *myopic* mechanism that maximizes instantaneous matching rate, we find that the optimal mechanism sets lower price to induce long-lived players to be pickier. This leads to higher matching quality, and also higher market thickness (e.g., more available long-lived players on the platform). Consistent with some prior studies, in our setting, the myopic mechanism is nearly optimal, namely, causing small social welfare loss. Interestingly, this implies that a central planner may use a range of nearly optimal prices to vary the trade-off between matching rate and matching quality, and change welfare distribution across the two sides of players.

In Section 5, with a homogeneous outside option of short-lived players, the solution of the optimal centralized mechanism also allows us to better compare the two optimal decentralized mechanisms. In particular, the long-lived-select mechanism indeed causes long-lived players to be pickier; as a result, the platform has to set a higher price to incentivize long-lived players to accept a matching angle closer to the optimal benchmark. However, such a high price yields zero surplus to the short-lived players, causing highly unbalanced welfare distribution across the two sides of players. In contrast, the short-lived-select mechanism can achieve the optimal social welfare and yield a much more balanced welfare allocation across the two sides.

In Section 6 with heterogeneous outside options of short-lived players, the platform's price affects short-lived players' participation. The price, together with short-lived players' arrival rate and degree of heterogeneity, determine the matching rate of short-lived players. When both the arrival rate of short-lived players and their heterogeneity are high, the short-lived-select mechanism can no longer achieve the centralized optimal social welfare, because the centralized optimal solution tends to set a higher price and a narrower matching threshold. In such a case, the status quo (i.e., long-lived-select mechanism) can outperform the short-lived-select mechanism in total social welfare. However, similar to findings in Section 5, if either the heterogeneity level or the arrival rate of short-lived players is low, the short-lived-select mechanism outperforms the status-quo by curbing the pickiness of long-lived players. Finally, in all cases, the heterogeneity protects short-lived players' surplus by discouraging the platform from setting too high a price; thus, there is much less concern of extremely unbalanced welfare distribution across the two sides of players, in comparison with the scenario of Section 5.

The rest of this paper is organized as follows. In Section 2, we compare and contrast results of our paper to existing literature. We introduce the model and formulate the matching system in Section 3. In Section 4, we introduce results when the platform design the matching threshold, and compare with results on decentralized systems in Section 5. We provide an extension involving heterogeneity among players in Section 6 and conclude our discussion in Section 7. All proofs and some detailed derivations are presented in the Appendix.

2. Literature

We highlight four streams of literature that are close to our paper. First, under centralized matching, we find that the platform only needs to design the price to induce desired matching thresholds. Furthermore, the optimal price, for which we have a closed-form approximation, is smaller than the myopic price which maximizes the matching rate. However, we also confirm that using the myopic price is still near optimal when players demonstrate low tolerance towards mismatch. A stream of related papers also study centralized matching and evaluate greedy/myopic matching policies under different settings. In particular, Anderson et al. (2017) consider a dynamic barter exchange system with different exchange rules, including: two- and three-way cycles, chain matching and their combinations. Different from our model, all players stay on the market until their goods are exchanged but incur waiting cost. They focus on minimizing the average waiting time and find that a greedy policy that conducts exchanges immediately as the opportunity arises is near optimal under all exchange rules they considered. Akbarpour et al. (2019) consider a similar setting to Anderson et al. (2017) with consideration of players' departure. They examine the benefit of knowing the exact departure time of each player in two-way exchanges. Furthermore, they show that if this information is available, letting players wait instead of matching greedily can benefit the platform significantly. Without this information, the greedy policy is near optimal. Ashlagi et al. (2019) study a dynamic matching market with easy and hard to match players. In their model, hard to match players have significantly lower matching probability compared with that of easy to match players, and all players want to be matched as soon as possible. They analyze the performance of myopic matching policies involving bilateral and chain matching. Different than these papers, who do not model the matching quality explicitly, we also study the trade-off between matching quality and matching rate.

Second, our closed-form results in Section 4 leverage on classic findings of $M/M/\infty$ queues, considering the number of long-lived players at any given time as the number of available servers. Many papers have also studied two-sided matching via queueing methods. Aféche et al. (2014) study trading systems using double-sided queues. They also consider short and long-lived players similar to our model without circular preference, and provide performance measures (such as expected waiting time, etc.) of queues under First-Come-First-Serve policy (FCFS). Büke and Chen (2017) consider two classes of players arriving at the system and each player can be matched with a player form the other class, just like our work. Both classes of players can stay on the market for exponentially distributed periods of time without matching. They analyze the fluid and diffusion limits of the matching system and the effect of the exogenous matching probability on system performance such as the average queue length of different classes. Many other papers also focus on system performance and stability conditions in matching systems for given policies, such as FCFS (see e.g., Caldentey et al. 2009, Adan and Weiss 2012). Furthermore, researchers have focused on analyzing different matching policies under fluid limits (see e.g., Zenios et al. 2000, Su and Zenios 2006, Akan et al. 2012, Gurvich and Ward 2014, Kanoria and Saban 2020). Recently, Özkan and Ward (2019) consider a matching problem for ride-sharing. They use players' origin or destination as types. Riders only accept drivers who can arrive in a certain time window. They take advantage of a large market, where players are considered as a continuum, and identify policies that match the most players. In our work, we do not scale the market size. Instead, we exploits the similarity between our matching system with $M/M/\infty$ queues when players demonstrate low tolerance on mismatch angles and obtain results under steady states.

Third, in Sections 5 and 6, we consider decentralized matching where long-lived players may behave strategically, which connects our paper to the literature on operational problems with strategic players. Chen and Hu (2020) study two-sided matching with forward-looking sellers and buyers. In their system, the platform decides both pricing and matching policies similar to our model. However, sellers and buyers decide when to participate in matching. That is, both sellers and buyers always monitor the dynamics of prices and decide when to send the platform requests to be matched. They show that a fixed price plus compensation for waiting costs together with greedy matching policy is asymptotically optimal when the market is large. Allon et al. (2012) analyze the role of a matching intermediary in a decentralized market. They compared different intervention methods and their impacts on the efficiency of decentralized matching. Arnosti et al. (2020) study a mean field game between applicants and employers in a large decentralized market. In their game, applicants decide their searching intensity while employers decides their screening strategies. Yang et al. (2016) use mean-field equilibrium to study the strategic relocation of drivers in ride-sharing. Liu et al. (2019) use data from DiDi to study a matching game empirically in a decentralized market of ride-sharing. They use the number of matches and the average matching quality to describe the efficiency of the system and demonstrate that increasing players' waiting time can improve the efficiency by increasing the market thickness. In our paper, we study the platform's decision on pricing and players' decisions on the acceptable matching quality, which is characterized by mismatch angles. In addition, we answer the question on which side of the market should decide the matching threshold for various settings.

Finally, our circular modeling approach resembles the Hotelling's circular city model in economics. It provides a tractable setup for compatibility of players' preferences that incorporates spatial features. The original model appears in Salop (1979), which uses a circular model as a geographical representation of a city. Suppliers and consumers in that model sit at fixed locations on a circle, and consumers have preferences over their relative locations to the suppliers. In our model, we have drivers and riders' destinations on the circle, and extend the circular city to include the center of the circle as the common origin. Recently, Pavan and Gomes (2019) extend circular city model to a three dimensional space as a cylinder, representing two dimensional preferences. However, there is no arrivals or departures in either of these two papers. Circular city model has also been applied in the Operations literature. Feng et al. (2020) consider a ride-hailing scenario in a circular city. In their model, riders arrive following a Poisson process with origins and destinations distributed on a circle. Furthermore, riders do not leave the system until being served. The number of taxis on the circle is fixed and they travel clock/counterclockwise with constant speed. Unlike our paper, they do not consider players' individual preferences. Their focus is on comparing the efficiency (average waiting time) between the traditional taxi services and that of the ride hailing services from a centralized perspective. In both of their mechanisms, a rider is always matched (or is picked up) with the nearest available taxi (or by the first taxi passing by) immediately.

3. Model Setup

Consider two classes $\{L, S\}$ of players arriving to the matching platform. Type L players are *long-lived* (patient). They follow a Poisson arrival process with rate λ_L. Each arrival has a "life time" that is exponentially distributed with rate γ; if no match occurs by the end of its life time, the player disappears from the platform at that time. Taking DiDi Hitch as an example, drivers may depart due to random events such as changes in traffic conditions, emergence of outside options, etc. Hence, their lifespan on the platform is random. Type S players are *short-lived* (impatient) with Poisson arrival rate λ_S, and leave the platform immediately if not matched upon arrival. Two players are matchable only if they are from different classes. Therefore, matches can only be made upon arrivals of short-lived players.

We use a "circular" model to describe compatibility between players. Upon arrival, each player from either class uniformly and independently claims a random spot on the edge of a circle. Between players from two classes, their mismatch is measured by the arc, or, equivalently, the central angel between their spots on the circle. To simplify notations in future sections, we define $\phi \in [0,1]$ as the *mismatch angle*, which is the ratio between the actual angle of two locations and the maximum possible angle π. For example, two players with angle $\pi/4$ (or 45 degrees) between their locations on the circle have a mismatch angle $\phi = 0.25$. In general, the mismatch angle ϕ is a value in $[0,1]$. Furthermore, in this paper, we assume that the long-lived player in each match bears the entire cost of mismatch, although our analytical method can also be readily extended to cases where players share (or split) the mismatch angle. In the carpooling example, a rider needs to be sent to her requested location and her driver needs to drive the entire detour. In the gig job example, the service needs to be provided exactly as requested by the service seeker, and the service provider needs to bear costs induced by the mismatch in skills such as buying additional tools and acquiring new knowledge, etc.

We assume that each short-lived player receives a benefit u per match while each long-lived player is subject to mismatch cost with coefficient c per unit of mismatch angle. We assume that u and c are fixed and observable by all parties and fixed in Sections 4 and 5. Taking DiDi Hitch as an example, this assumption reflects areas of a city where short-lived players' alternative travel options are the same and long-lived players are facing the same operating condition. Furthermore, the platform sets a payment P from a short-lived player to a long-lived player in each successful match. Therefore, we define short- and long-lived player's utilities in a match with mismatch angel $\phi \in [0,1]$ and payment $P \in [0,u]$ as,

$$W_S(\phi) = u - P, \text{ and } W_L(\phi) = P - c\phi, \text{ respectively.} \tag{3.1}$$

As we can see from (3.1), short-lived players' utility is fixed in each match despite the quality of the match as they do not bear any mismatch angle. However, long-lived players' utility function is decreasing *w.r.t.* the mismatch angle. We restrict $P \leq u$ to ensure short-lived players' participation.

Following the utility functions in (3.1), a long-lived player shall participate in a match only if the mismatch angle is no greater than P/c. Therefore, define mismatch *tolerance* as

$$\epsilon := \frac{u}{c}. \tag{3.2}$$

The larger ϵ is, the more tolerant players are towards mismatch angles. Furthermore, define normalized *price* as

$$\rho := \frac{P}{u}, \tag{3.3}$$

representing the fraction of a short-lived player's benefit transfered to a long-lived player in each match. For example, if the price $\rho = 1$, the platform leaves a short-lived player 0 surplus. The smaller ρ is, the more benefit a short-lived player can keep from a match.

We define the social welfare generated in each match as the summation of utilities from both sides. In other words, given a mismatch angle $\phi \in [0,1]$ in a match, the social welfare is

$$W_{SW}(\phi) = u - c\phi = c(\epsilon - \phi). \tag{3.4}$$

Since the transfer between players are internal to the system, the social welfare in (3.4) is simply the difference between the benefit generated for a short-lived player and the mismatch penalty to a long-lived player.

Throughout this paper, we consider the platform always announces its price ρ first, anticipating players' behaviors. In Section 4, we consider centralized matching that the platform also decides a matching threshold to maximize social welfare. For tractability, we assume the threshold does not depend on any information other than the price ρ, so that we denote $\Theta(\rho)$ as the matching threshold. To be more specific, whenever a long-lived player has a mismatch angle no greater than $\Theta(\rho)$ with an arriving short-lived player, they are matched immediately by the platform. In Section 5, we consider two decentralized matching systems: *long-lived-select* and *short-lived-select* matching. Under long-lived-select matching, long-lived players can choose whether to match with an arriving short-lived player or not. The platform does not reveal any other information to each long-lived player during his lifespan on the platform, which is exponentially distributed. Thus, long-lived-select matching is equivalent to a direct mechanism where each long-lived player reports a matching threshold $\Theta(\rho)$ to the platform and shall be matched with any short-lived player who first appears within the threshold. Under short-lived-select matching, upon arrival, each short-lived player can choose to match with any long-lived player who earns a non-negative utility from matching. In all three matching mechanisms above, as a tie-breaking rule, the pair of players who has the least mismatch angle shall be matched, if there are multiple pairs satisfying the matching criteria.

For the rest of this section, we first characterize the dynamics of long-lived players over time and then define the platform's objective.

3.1. Matching probability and Birth-Death process.

Consider any matching threshold $\Theta(\rho) = \theta$, such that $\theta \in [0, 1]$. Suppose upon the arrival of a short-lived player, x long-lived players are available. Denote ϕ_i, $i \in \{1, ..., x\}$ as the mismatch angle between each of the x long-lived players and the focal short-lived player. Let $\underline{\phi} = \min\{\phi_i \,|\, i = 1, ..., x\}$ represent the minimum mismatch angle between the short-lived player and the x long-lived players. According to the threshold matching policy, a match can be made only if $\underline{\phi} \leq \theta$. Thus, the matching probability under threshold $\theta \in [0, 1]$ is

$$p_x(\theta) = 1 - \Pi_{i=1}^{x} \mathbb{P}(\phi_i > \theta) = 1 - (1 - \theta)^x, \tag{3.5}$$

where $1 - \theta$ is the probability that the short-lived player cannot be matched with a long-lived player, whose location is uniformly distributed on the circle. Since long-lived players are located independently, $(1 - \theta)^x$ is the probability that all x number of long-lived players cannot be matched with the short-lived player.

Since the short-lived player can only be matched to the long-lived player with the minimum mismatch angle, it is important to characterize the distribution function of $\underline{\phi}$. Note that the probability that the minimum mismatch angle $\underline{\phi}$ is no greater than a value $\phi \in [0, 1]$ is simply $p_x(\phi)$

as defined in (3.5). By differentiating function $p_x(\phi)$ with respect to ϕ, we obtain the probability density function of the minimum mismatch angle $\underline{\phi}$ as

$$g_x(\underline{\phi}) = x(1-\underline{\phi})^{x-1}, \quad \forall\, 0 \leq \underline{\phi} \leq 1. \tag{3.6}$$

Since we focus on threshold policies on the minimum mismatch angle, the distribution of $\underline{\phi}$ helps us characterize the quality of matching outcomes in later sections.

As we can see, both matching probability and distribution of the minimum mismatch angle depend on the number x of long-lived players on the platform. To characterize the dynamics of the long-lived players, we formulate the arrival and departure of long-lived players as a Continuous-Time Markov Chain (CTMC), which is a Birth-Death process. Denote $f_\theta(x)$ to represent the probability mass function of the stationary distribution of this Birth-Death process if the matching threshold is θ. That is, $f_\theta(x)$ is the steady state probability that there are x long-lived players on the market, which solves the following system of equations,

$$\lambda_L f_\theta(0) = (\lambda_s p_0(\theta) + \gamma) f_\theta(1),$$

$$\lambda_L f_\theta(x-1) + (\lambda_s p_{x+1}(\theta) + \gamma(x+1)) f_\theta(x+1) = (\lambda_L + \lambda_S p_x(\theta) + \gamma x) f_\theta(x), \quad x = 1, 2, \ldots$$

This system of difference equations has the following unique solution[1]:

$$f_\theta(x) = \frac{{\lambda_L}^x f_\theta(0)}{\Pi_{k=1}^x (\lambda_S p_k(\theta) + \gamma k)}, \quad \forall x \geq 1, \tag{3.7}$$

and

$$f_\theta(0) = \left(1 + \sum_{x=1}^{\infty} \frac{{\lambda_L}^x}{\Pi_{k=1}^x (\lambda_S p_k(\theta) + \gamma k)}\right)^{-1}. \tag{3.8}$$

Consider any matching threshold $\Theta(\rho)$ that only depends on the price ρ. We denote $X(\Theta(\rho))$ as the random variable according to the steady state distribution $f_{\Theta(\rho)}(\cdot)$. In the following sections, we conduct analysis under steady state of the matching system where the number of long-lived players follows the distribution $f_{\Theta(\rho)}(\cdot)$. As a result, all utility functions are evaluated by taking expectations over the random variable $X(\Theta(\rho))$.

3.2. Social welfare and other performance measures.

We consider social welfare maximization as the platform's primary objective since it reflects both matching quality and quantity. Throughout this paper, unless specified otherwise, we refer to the expected social welfare generated per arrival of short-lived player as the *social welfare*, since matching can only happen upon their arrival and the arrival rate is a constant. Thus, define the expected social welfare as

$$U(\theta) := \mathbb{E}_{X(\theta)} \mathbb{E}_{\underline{\phi}} \left[W_{SW}(\underline{\phi}) \mathbb{I}\{\underline{\phi} \leq \theta\} \mid X(\theta) \right] = \mathbb{E}_{X(\theta)} \mathbb{E}_{\underline{\phi}} \left[c(\epsilon - \underline{\phi}) \mathbb{I}\{\underline{\phi} \leq \theta\} \mid X(\theta) \right], \tag{3.9}$$

where function $W_{SW}(\cdot)$ is the social welfare generated per match as defined in (3.4), and θ represents the matching threshold. In (3.9), the inner expectation $\mathbb{E}_{\underline{\phi}}\left[W_{SW}(\underline{\phi})\mathbb{I}\{\underline{\phi} \le \theta\} \,|\, X(\theta)\right]$ represents the expected social welfare generated in a match given the number of long-lived players on the platform. It is taken with respect to the random variable $\underline{\phi}$, representing the minimum mismatch angle, which follows distribution function (3.6). The outer expectation captures the random number of long-lived players that an arriving short-lived player sees, following P.A.S.T.A (Possion Arrival See Time Average). That is, the number of long-lived players follows the same steady state distribution upon the arrival of each short-lived player (see, e.g., Wolff 1982). Note that function U in (3.9) represents the expected social welfare generated per arrival of short-lived player under steady state. The social welfare rate is $\lambda_S U(\theta)$ per unit of time because matching can only happen upon short-lived players' arrivals.

Because maximizing social welfare is the platform's primary objective, in Section 4, the platform designs both the price ρ and the matching threshold $\Theta(\rho)$ to maximize $U(\Theta(\rho))$. In Section 5, the platform still controls ρ but the threshold $\Theta(\rho)$ is an equilibrium outcome chosen by either short- or long-lived players.

Besides the primary objective of maximizing social welfare, the platform may also consider other secondary performance measures such as matching rate and welfare allocation. Under steady state with matching threshold θ, the matching rate per unit of time is

$$R(\theta) := \lambda_S \mathbb{E}_{X(\theta)}\left[p_{X(\theta)}(\theta)\right] = \lambda_S \mathbb{E}_{X(\theta)}\left[1 - (1-\theta)^{X(\theta)}\right], \tag{3.10}$$

where function $p_x(\cdot)$, representing the matching probability, is defined in (3.5). Since matching can only happen upon the arrival of an short-lived player, the matching rate is the product of short-lived players' arrival rate and the matching probability. As for welfare allocation, with price ρ and threshold θ, the expected utility of a short-lived player is

$$S(\rho,\theta) := \mathbb{E}_{X(\theta)}\left[\mathbb{E}_{\underline{\phi}}\left[W_S(\underline{\phi})\mathbb{I}\{\underline{\phi} \le \theta\} \,|\, X(\theta)\right]\right] = \mathbb{E}_{X(\theta)}\left[\mathbb{E}_{\underline{\phi}}\left[\epsilon c(1-\rho)\mathbb{I}\{\underline{\phi} \le \theta\} \,|\, X(\theta)\right]\right], \tag{3.11}$$

where function W_S is defined in (3.1). Note that welfare generated by long-lived players has rate $\lambda_S \mathbb{E}_{X(\theta)}\mathbb{E}_{\underline{\phi}}\left[W_L(\underline{\phi})\mathbb{I}\{\underline{\phi} \le \theta\} \,|\, X(\theta)\right]$. Thus, the expected utility of an individual long-lived player is

$$L(\rho,\theta) := \frac{\lambda_S}{\lambda_L}\mathbb{E}_{X(\theta)}\left[\mathbb{E}_{\underline{\phi}}\left[W_L(\underline{\phi})\mathbb{I}\{\underline{\phi} \le \theta\} \,|\, X(\theta)\right]\right] = \frac{\lambda_S}{\lambda_L}\mathbb{E}_{X(\theta)}\left[\mathbb{E}_{\underline{\phi}}\left[c(\epsilon\rho - \underline{\phi})\mathbb{I}\{\underline{\phi} \le \theta\} \,|\, X(\theta)\right]\right], \tag{3.12}$$

where λ_L is the arrival rate of long-lived players. Note that we have $\lambda_S S(\rho,\theta) + \lambda_L L(\rho,\theta) = \lambda_S U(\theta)$, representing the social welfare is made by the summation of short- and long-lived players' utilities of matching.

4. Centralized Matching

In this section, the platform chooses price $\rho \in [0,1]$ and the matching threshold $\Theta(\rho)$ to maximize social welfare. Without loss of generality, we can focus on $\Theta(\rho) \in [0, \epsilon\rho]$ since $\Theta(\rho) > \epsilon\rho$ generates negative utilities for long-lived players. This matching mechanism is easy to implement in reality. Taking DiDi Hitch as an example, the platform only needs to announce a fixed normalized price ρ that each rider (short-lived) needs to pay to a driver (long-lived). Upon the arrival of a rider, the platform matches her with a driver, with whom, she has the smallest mismatch angle that is no greater than $\Theta(\rho)$. To ensure drivers' acceptance, the platform commits to exclude anyone who has declined a matchable rider from future matching. As long as $\Theta(\rho) \le \epsilon\rho$, it is in the best interests of drivers to accept matches purposed by the platform.

Note that each price ρ induces a maximum mismatch angle $\epsilon\rho$ that a long-lived player has a non-negative utility when matched. Thus, when the platform is in charge of both ρ and $\Theta(\rho)$, it is optimal to set

$$\Theta(\rho) = \epsilon\rho.$$

The reason is that price is internal to the social welfare function according to (3.9) and therefore, using a threshold $\Theta(\rho) < \epsilon\rho$ is equivalent to setting a price $\Theta(\rho)/\epsilon$, which is smaller than ρ. Thus, the platform's problem in (3.9) is equivalent to

$$\max_{\rho \in [0,1]} U(\epsilon\rho) = \max_{\rho \in [0,1]} \mathbb{E} h_\epsilon(X(\epsilon\rho), \epsilon\rho), \tag{4.1}$$

where

$$h_\epsilon(x, \theta) := \mathbb{E}_{\underline{\phi}} \left[W_{SW}(\underline{\phi}) \mathbb{I}\{\underline{\phi} \le \theta\} \,|\, x \right] = \mathbb{E}_{\underline{\phi}} \left[c(\epsilon - \underline{\phi}) \mathbb{I}\{\underline{\phi} \le \theta\} \,|\, x \right]. \tag{4.2}$$

Function $h_\epsilon(x, \theta)$ represents the expected social welfare generated in a match under matching threshold θ when there are x number of long-lived players on the platform. Also note that when the matching threshold equals to $\epsilon\rho$, the matching rate in (3.10) becomes

$$R(\epsilon\rho) = \lambda_S \mathbb{E}_{X(\epsilon\rho)}[1 - (1 - \epsilon\rho)^{X(\epsilon\rho)}], \quad \forall \rho \in [0,1]. \tag{4.3}$$

It is intuitive that the matching rate $R(\epsilon\rho)$ is increasing in ρ. We provide a formal proof of this claim in Appendix B.2. Thus, we denote $\rho_M := 1$ as the *myopic price*, which maximizes the matching rate, serving as a benchmark for our analysis.

It is hard to obtain analytical results on the optimal price for the matching system proposed so far. Even though (4.2) has a closed-form expression, its expectation with respect to $X(\epsilon\rho)$ is hard to compute in closed-form due to complicated expressions of its steady state distribution. Therefore, we propose an approximation approach in this section when ϵ approaches 0.

We first consider an $M/M/\infty$ queue related to our original CTMC, with arrival rate $\lambda = \lambda_L$ and service rate $\mu = \gamma + \lambda_S \epsilon\rho$. Note that this $M/M/\infty$ queue mimics the original Birth-Death process. The x active servers at any point in time resemble x long-lived players on the platform. The matching rate corresponds to system departure rate $x(\epsilon\rho+\gamma)$. Thus, the matching probability in this system is simply $\epsilon\rho x$, when there are x long-lived players. The next lemma shows the relationship between matching probabilities under the original process and the $M/M/\infty$ queue.

LEMMA 1. *Fix ϵ and ρ such that $0 \le \epsilon\rho \le 1$, and consider the matching probability function $p_x(\cdot)$ in (3.5). We have*

$$\epsilon\rho x \ge p_x(\epsilon\rho) = 1 - (1-\epsilon\rho)^x, \quad \forall x \ge 1.$$

The intuition behind this upper bound of matching probability is straightforward. Note that $\epsilon\rho x$ is the matching probability when the x long-lived players' tolerance angles do not overlap on the circle. Thus, their total tolerance covers the maximum possible angle on the circle, which leads to a greater matching probability. Therefore, $\epsilon\rho x$ is an upper bound of $p_x(\epsilon\rho)$, and this $M/M/\infty$ system's departure rate is higher than or equal to the one in the original process.

Following classic results on $M/M/\infty$ queues (see, e.g. Iglehart 1965), the stationary distribution of the aforementioned $M/M/\infty$ queue is a Poisson random variable, with parameter $\dfrac{\lambda_L}{\lambda_S\epsilon\rho+\gamma}$. Denote $Y(\epsilon\rho)$ to represent this Poisson random variable. Similarly, we can consider another Poisson random variable $Y(0)$ to be the stationary distribution of an $M/M/\infty$ queue with departure rate γ (that is, $\epsilon = 0$ in $\lambda_S\epsilon\rho+\gamma$), which is stochastically no greater than the one in the Birth-Death process for all states. The next proposition formalizes the relationship between $Y(0)$, $Y(\epsilon\rho)$ and $X(\epsilon\rho)$.

PROPOSITION 1. *Fixing $\epsilon \in [0,1]$ and $\rho \in [0,1]$, we have*

$$Y(0) \succeq_1 X(\epsilon\rho) \succeq_1 Y(\epsilon\rho). \tag{4.4}$$

Moreover, we have

$$\mathbb{E}h_\epsilon(Y(\epsilon\rho),\epsilon\rho) \le \mathbb{E}h_\epsilon(X(\epsilon\rho),\epsilon\rho) \le \mathbb{E}h_\epsilon(Y(0),\epsilon\rho). \tag{4.5}$$

We use the notation $\succeq_1$ to represent the first order stochastic dominance between two random variables[2]. Proposition 1 shows that replacing $X(\epsilon\rho)$ with $Y(\epsilon\rho)$ and $Y(0)$ lead to a *lower* and an *upper* bound for the platform's objective function in (4.1), respectively.

From this point forwards, we use $Y(\epsilon\rho)$ to replace $X(\epsilon\rho)$ and let ϵ approach 0 when seeking analytical results. The next proposition justifies this approximation.

PROPOSITION 2. *Consider function h_ϵ in (4.2) and fix $\rho \in [0,1]$. We have*

$$\lim_{\epsilon\to 0}\frac{1}{\epsilon}\left(\frac{\mathbb{E}h_\epsilon(X(\epsilon\rho),\epsilon\rho)-\mathbb{E}h_\epsilon(Y(\epsilon\rho),\epsilon\rho)}{\mathbb{E}h_\epsilon(X(\epsilon\rho),\epsilon\rho)}\right)\le \lim_{\epsilon\to 0}\frac{1}{\epsilon}\left(\frac{\mathbb{E}h_\epsilon(Y(0),\epsilon\rho)-\mathbb{E}h_\epsilon(Y(\epsilon\rho),\epsilon\rho)}{\mathbb{E}h_\epsilon(Y(\epsilon\rho),\epsilon\rho)}\right)=\frac{\lambda_S\rho}{\gamma}. \quad (4.6)$$

In (4.6), the relative difference $\frac{\mathbb{E}h_\epsilon(X(\epsilon\rho),\epsilon\rho)-\mathbb{E}h_\epsilon(Y(\epsilon\rho),\epsilon\rho)}{\mathbb{E}h_\epsilon(X(\epsilon\rho),\epsilon\rho)}$ captures the quality of the approximation when using the Poisson random variable $Y(\epsilon\rho)$ in the objective function. Note that both the numerator and the denominator in the relative difference converge to 0 when ϵ approaches 0. With the help of Proposition 1, we have the first inequality in (4.6). Proposition 2 further states that the relative difference converge to 0 as ϵ approaches 0. Therefore, using random variable $Y(\epsilon\rho)$ instead of $X(\epsilon\rho)$ in the objective function is a good approximation when ϵ approaches 0. Furthermore, Proposition 2 sheds light on searching for analytical results for the social welfare rate in (4.1), because both $\mathbb{E}h_\epsilon(Y(\epsilon\rho),\epsilon\rho)$ and $\mathbb{E}h_\epsilon(Y(0),\epsilon\rho)$ have closed-form expressions. In this paper, we focus on using the Poisson random variable $Y(\epsilon\rho)$ instead of $Y(0)$. Although $Y(0)$ is also a good approximation for $X(\epsilon\rho)$ when ϵ goes to 0, such an approximation effectively assumes matching has no impact on the stationary distribution of the number of long-lived players. In other words, the random variable $Y(0)$ implies a matching system with replacement. That is, whenever a match is made, the system replaces the matched long-lived player with a new one at the same position on the circle until the exponential random time runs out. Such a matching system promotes aggressive matching (with higher price and larger matching thresholds). In fact, one can verify that using random variable $Y(0)$ in the objective function leads to the myopic price $\rho = 1$.

The expression $\mathbb{E}h_\epsilon(Y(\epsilon\rho),\epsilon\rho)$ (provided in the appendix) is still complex. Therefore, we perform a Taylor expansion for $\mathbb{E}h_\epsilon(Y(\epsilon\rho),\epsilon\rho)$ around $\epsilon = 0$. The next lemma provides a simple approximation when letting ϵ approach 0.

LEMMA 2. *Consider the Poisson random variable $Y(\epsilon\rho)$ with rate parameter $\frac{\lambda_L}{\lambda_S\epsilon\rho+\gamma}$. There exists a third order polynomial J of ρ and ϵ, such that,*

$$\lim_{\epsilon\to 0}\frac{1}{\epsilon^3}\left|\mathbb{E}h_\epsilon(Y(\epsilon\rho),\epsilon\rho)-J(\rho,\epsilon)\right|=0,\ \ \textit{or},\ \ \mathbb{E}h_\epsilon(Y(\epsilon\rho),\epsilon\rho)=J(\rho,\epsilon)+o(\epsilon^3). \quad (4.7)$$

We provide the exact expression of function $J(\rho,\epsilon)$ in the proof in the appendix. Here $o(\epsilon^k)$ represents terms such that $\lim_{\epsilon\to 0}\frac{o(\epsilon^k)}{\epsilon^k}=0$, for a value $k\ge 0$. That is, these terms approach 0 faster than ϵ^k.

Lemma 2 greatly simplifies expressions of social welfare rate in (4.1) under the limiting regime when ϵ goes to 0. In particular, the expected utility can be approximated by $J(\rho,\epsilon)$, which is a simple polynomial of ρ. We can then maximize function $J(\rho,\epsilon)$ over ρ towards an approximation of the optimal price. Proposition 2 and Lemma 2 are the building blocks for approximately solving platform's problems (4.1) in closed form. The following theorem summarizes our result.

Theorem 1. *Denote $\rho^* \in [0,1]$ as the optimal solution to (4.1), and $\rho^*_\epsilon \in [0,1]$ to be a maximizer of $J(\rho,\epsilon)$ for a given ϵ. Consider $\epsilon < \dfrac{\gamma}{\lambda_s + \lambda_L}$, we have*

$$\rho^*_\epsilon = \hat{\rho}_\epsilon + o(\epsilon), \tag{4.8}$$

in which

$$\hat{\rho}_\epsilon = 1 - \frac{\lambda_S}{2\gamma}\epsilon. \tag{4.9}$$

Furthermore, we have,

$$\lim_{\epsilon \to 0} \frac{1}{\epsilon}\left(\frac{U(\epsilon\rho^*) - U(\epsilon\hat{\rho}_\epsilon)}{U(\epsilon\rho^*)}\right) \leq \frac{\lambda_S}{\gamma}. \tag{4.10}$$

where function social welfare U is defined in (3.9).

In Theorem 1, (4.9) provides an approximation, $\hat{\rho}_\epsilon$, to the price that maximizes function J. The expression of $\hat{\rho}_\epsilon$ reveals further insights. We observe that $\hat{\rho}_\epsilon$ is always no greater than the myopic price $\rho_M = 1$, which extracts all surplus from short-lived players and induces the largest matching rate. Note that a smaller price (that yields a smaller matching threshold) means lower matching rate but better quality in each match on average. Therefore, the platform uses a lower price to trade-off between matching quantity and quality. This is consistent with strategic delay/waiting strategies identified in the matching literature (see e.g. Akbarpour et al. 2019, Liu et al. 2019).

Furthermore, the expression of the optimal centralized price in (4.9) shows the fundamental connection between *market thickness* and platform's matching decisions. In particular, we argue that the thicker the market is, the smaller the price $\hat{\rho}_\epsilon$ is. We can interpret the ratio $\dfrac{\lambda_S}{\gamma}$ as the average number of short-lived players that a long-lived player shall encounter during the time on the platform. Intuitively, the larger the value of λ_S is, the thicker the short-lived players' side of the market is. As a result, each long-lived player encounters more potential matching partners on average, which leads to better matching quality. Similarly, as γ decreases, long-lived players stay for a longer period of time on average, which leads to a thicker market as well. Furthermore, we observe that the gap between the optimal and myopic thresholds disappears when γ goes to infinity, representing the case where long-lived customers become more like short-lived. In this case, the platform cannot do better than matching myopically.

Interestingly, there is a secondary effect of using a smaller matching threshold that decreases the matching rate. The platform can intentionally build market thickness by being pickier than matching myopically. We can see the effect of using a smaller price by looking at the Poisson random variable $Y(\epsilon\rho)$, which represents the number of long-lived players on the platform. Note that the mean of $Y(\epsilon\rho)$ is $\dfrac{\lambda_L}{\lambda_S\epsilon\rho + \gamma}$, which increases as ρ decreases. The same result holds for the random variable $X(\epsilon\rho)$ as well. In other words, the platform gives up some marginal welfare from matches with large mismatch angle to build a thicker market of long-lived players. However, under a thicker market, the benefit of better matching quality outweighs the slight reduction in matching rate from using a lower price.

In (4.10) of Theorem 1, the term $\dfrac{U(\epsilon\rho^*)-U(\epsilon\hat{\rho}_\epsilon)}{U(\epsilon\rho^*)}$ represents the relative sub-optimality when using our approximation with price $\hat{\rho}_\epsilon$. It captures the quality of using the approximation $\hat{\rho}_\epsilon$ in the original objective function. Similar to the relative difference in Proposition 2, both the numerator and denominator of the relative sub-optimality converge 0 when ϵ approaches to 0. We bound this relative sub-optimality using Proposition 1. Thus, (4.10) confirms that the relative sub-optimality converges to 0 together with ϵ. Therefore, the approximate optimal price $\hat{\rho}_\epsilon$ performs well when ϵ is small. In fact, for any price that is a linear combination of the approximate optimal price $\hat{\rho}_\epsilon$ and the myopic price ρ_M, we can show that it induces near optimal social welfare when ϵ approaches 0. We formalize this result in the next corollary.

Corollary 1. *Consider price*

$$\rho_\epsilon(\alpha)=\alpha\rho_M+(1-\alpha)\hat{\rho}_\epsilon=1-(1-\alpha)\frac{\lambda_S}{2\gamma}\epsilon, \quad \forall\,\alpha\in[0,1]. \tag{4.11}$$

We have

$$\lim_{\epsilon\to 0}\frac{1}{\epsilon}\left(\frac{U(\epsilon\rho^*)-U(\epsilon\rho_\epsilon(\alpha))}{U(\epsilon\rho^*)}\right)\le\frac{\lambda_S}{2\gamma}, \quad \forall\,\alpha\in[0,1]. \tag{4.12}$$

Corollary 1 states that price $\rho_\epsilon(\alpha)$, which is in between the approximate optimal price and the myopic price, induces near optimal social welfare in our limiting regime. It further implies that the myopic price $\rho_M=1$ is near optimal ($\alpha=1$). Thus, the platform can use a variety of prices according to (4.11) that leads to different combinations of matching quantity and quality while achieving near optimal social welfare.

We end this section by providing a numerical example that shows the welfare allocation among players by computing individual utility functions S and L defined in (3.11) and (3.12), respectively. In Figure 1, we plot the social welfare $U(\epsilon\rho_\epsilon(\alpha))$, short-lived players' individual utility $S(\rho_\epsilon(\alpha),\epsilon\rho_\epsilon(\alpha))$ and long-lived players' individual utility $L(\rho_\epsilon(\alpha),\epsilon\rho_\epsilon(\alpha))$, with respect to different prices with values of $\alpha\in[0,1]$. Note that matching myopically ($\alpha=1$) takes away the entire surplus of short-lived players, which leads to significant unbalance in welfare distribution among players. Furthermore, although decreasing the value of α, i.e., moving from using the myopic price to the approximated optimal price, leads to minor improvement in overall social welfare, it may drastically improve the balance in welfare allocation among players.

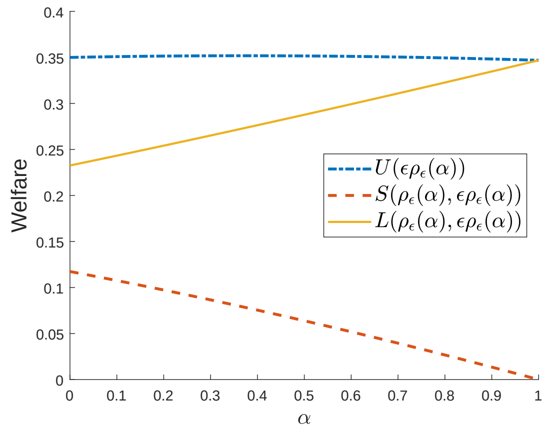

Figure 1 Welfare allocations with respect to different values of α, with parameters $\gamma = 1.2$, $\lambda_L = \lambda_S = 10$, and $c = 50$

5. Decentralized Matching

In this section, we study situations under which the platform sets the price and lets either side to decide the matching threshold $\Theta(\rho)$. That is, the platform designs the price ρ to maximize social welfare (3.9), anticipating players' behaviors.

We first focus on the long-lived-select matching that reflects the status quo of the current matching mechanism in DiDi Hitch. In this case, the platform allows long-lived players to select each arriving short-lived player. However, the platform does not reveal any other information to them, such as the number or locations of other long-lived players. Since each long-lived player's lifespan is exponentially distributed, they use a static matching threshold $\Theta(\rho)$ that depends on the only information, price ρ. Thus, long-lived-select matching is equivalent to letting each long-lived player commit to a matching threshold $\Theta(\rho)$ prior to joining the platform. By reporting his matching threshold $\Theta(\rho)$ to the platform, a long-lived player allows the platform to match him with a short-lived player if the mismatch angle is no greater then $\Theta(\rho)$ and is the minimum mismatch angle among all long-lived players.

When choosing the threshold $\Theta(\rho)$, a focal long-lived player needs to anticipate other long-lived players' thresholds, which affect his matching outcome. So it is a game with potentially infinitely many long-lived players. We focus on symmetric equilibria in the long-lived players' game. Denote $\theta(\rho)$ as a long-lived players' equilibrium matching threshold. That is, in equilibrium, all long-lived players shall choose matching threshold $\Theta(\rho) = \theta(\rho)$ such that no long-lived player shall unilaterally deviate. Because each player's threshold only depends on the price ρ and does not change over time, the equilibrium threshold also only depends on the price ρ, hence, the notation $\theta(\rho)$. Therefore, we can rewrite the platform's problem in (3.9) as

$$\max_{0\le\rho\le 1} U(\theta(\rho)), \tag{5.1}$$

with equilibrium threshold $\theta(\rho)$ induced by price ρ.

In this section, we do not consider the limiting regime where ϵ approaches 0, because it does not simplify our analysis. We first derive the expected utility function for each long-lived player and provide a heuristic method to evaluate it efficiently. Next, we identify the equilibrium matching threshold $\theta(\rho)$ numerically. Then, we calculate the platform's optimal price when long-lived players choose their matching thresholds ρ_L, which optimizes (5.1), and numerically compare it with the exact centralized optimal price ρ^*, which optimizes (4.1).

To find the equilibrium matching threshold $\theta(\rho)$ induced by price ρ, consider the situation where all other x long-lived players commit to a threshold θ and the focal player uses a threshold $\hat{\theta}$. Without loss of generality, we only consider $\theta,\hat{\theta}\in[0,\epsilon\rho]$, which guarantee that utilities are always non-negative.

Note there are two immediate outcomes upon an arrival of a short-lived player. First, the focal player may be matched with a short-lived player. Following the distribution function of minimum mismatch angle (3.6), we define function

$$\mathcal{A}(x,\hat{\theta},\theta,\rho) := \int_0^{\theta}\int_0^{\phi} c(\epsilon\rho-\hat{\phi})d\hat{\phi}g_x(\phi)d\phi + \int_{\theta}^{1}\int_0^{\hat{\theta}} c(\epsilon\rho-\hat{\phi})d\hat{\phi}g_x(\phi)d\phi, \quad \forall\theta,\hat{\theta}\in[0,\epsilon\rho] \text{ and } x\ge 0. \tag{5.2}$$

Then we have the following expression for a long-lived player's expected utility,

$$\mathbb{E}_{\hat{\phi},\underline{\phi}}\left[W_L(\hat{\phi})\mathbb{I}\{\hat{\phi}\le\min\{\underline{\phi},\hat{\theta}\}\}\,|\,x\right] = \begin{cases}\mathcal{A}(x,\hat{\theta},\hat{\theta},\rho), & \text{if } \hat{\theta}\le\theta,\\ \mathcal{A}(x,\hat{\theta},\theta,\rho), & \text{if } \hat{\theta}>\theta,\end{cases} \tag{5.3}$$

where function W_L is defined in (3.1); random variable $\hat{\phi}$ represents the mismatch angle between the focal player and a short-lived player, whose position follows a uniform distribution; random variable $\underline{\phi}$ represents the minimum mismatch angle among all other players, following the distribution in (3.6); and the indicator function inside the expectation states the condition that the focal player shall be matched. Together, the expression in (5.3) represents the focal player's expected utility upon the time a short-lived player arrives, when she commits to a threshold $\hat{\theta}$, and faces x *other* long-lived players who use threshold θ.

Second, the focal player may be unmatched when competing with the other x long-lived players. Define function

$$\mathcal{B}(x,\hat{\theta},\theta) = \int_0^{\hat{\theta}}\int_{\phi}^{1} d\hat{\phi}h_x(\phi)d\phi + \int_{\hat{\theta}}^{\theta}\int_{\hat{\theta}}^{1} d\hat{\phi}h_x(\phi)d\phi, \quad \forall\theta,\hat{\theta}\in[0,\epsilon\rho] \text{ and } x\ge 1. \tag{5.4}$$

Then we have

$$\mathbb{P}\left(\hat{\phi} > \min\{\underline{\phi}, \hat{\theta}\} \text{ and } \underline{\phi} \le \theta \,\middle|\, x\right) = \begin{cases} \mathcal{B}(x, \hat{\theta}, \theta), & \text{if } \hat{\theta} \le \theta, \\ \mathcal{B}(x, \theta, \theta), & \text{if } \hat{\theta} > \theta, \end{cases} \tag{5.5}$$

representing the probability that the focal player with threshold $\hat{\theta}$ is unmatched while one of the other x long-lived player with threshold θ is matched. Furthermore, define function

$$\mathcal{C}(x, \hat{\theta}, \theta) = \int_{\theta}^{1} \int_{\hat{\theta}}^{1} d\hat{\phi} h_x(\phi) d\phi, \quad \forall \theta, \hat{\theta} \in [0, \epsilon\rho] \text{ and } x \ge 1, \tag{5.6}$$

and we have

$$\mathbb{P}\left(\hat{\phi} > \min\{\underline{\phi}, \hat{\theta}\} \text{ and } \underline{\phi} > \theta \,\middle|\, x\right) = \mathcal{C}(x, \hat{\theta}, \hat{\theta}), \tag{5.7}$$

representing the probability that no player is matched.

Now we define $V(x, \hat{\theta}, \theta, \rho)$ as the expected utility function for a focal player when the platform's price is ρ, all other x long-lived players commit to a matching threshold θ, and the focal one uses a threshold $\hat{\theta}$. Dropping variables $\hat{\theta}, \theta$ and ρ in function V when there is no confusion, a function $V(x)$ solves according to the following difference equation for $x \ge 1$:

$$V(x) = \begin{cases} \dfrac{\lambda_S \mathcal{A}(x, \hat{\theta}, \hat{\theta}, \rho) + \left[\lambda_S \mathcal{B}(x, \hat{\theta}, \theta) + x\gamma\right] V(x-1) + \lambda_L V(x+1)}{\lambda_L + \lambda_S(1 - \mathcal{C}(x, \hat{\theta}, \theta)) + (x+1)\gamma}, & \text{if } \hat{\theta} \le \theta, \\ \dfrac{\lambda_S \mathcal{A}(x, \hat{\theta}, \theta, \rho) + \left[\lambda_S \mathcal{B}(x, \theta, \theta) + x\gamma\right] V(x-1) + \lambda_L V(x+1)}{\lambda_L + \lambda_S(1 - \mathcal{C}(x, \hat{\theta}, \theta)) + (x+1)\gamma}, & \text{if } \hat{\theta} > \theta, \end{cases} \tag{5.8}$$

with boundary condition

$$\lambda_S \mathcal{A}(0, \hat{\theta}, \theta, \rho) = (\lambda_L + \lambda_S \mathcal{C}(0, \hat{\theta}, \theta) + \gamma) V(0) - \lambda_L V(1), \tag{5.9}$$

where functions $\mathcal{A}$, $\mathcal{B}$ and $\mathcal{C}$ are defined in (5.2), (5.4) and (5.6), respectively. We provide the derivation of the difference equations in Appendix B.3.

Assume that upon each long-lived player's arrival, the matching system already reaches its steady state. Define a long-lived player's expected utility upon arriving at the platform if he uses threshold $\hat{\theta}$ while all other long-lived players use θ as

$$\mathbb{E}[V(X(\theta), \hat{\theta}, \theta, \rho)], \tag{5.10}$$

where the expectation follows P.A.S.T.A. Furthermore, we formally define the long-lived players' equilibrium matching threshold as

$$\theta(\rho) \in \arg\max_{\hat{\theta} \in [0, \epsilon\rho]} \mathbb{E}[V(X(\theta(\rho)), \hat{\theta}, \theta(\rho), \rho)]. \tag{5.11}$$

Unfortunately, we do not have a theoretical guarantee for the existence or the uniqueness of the equilibrium matching threshold because the game lacks desirable structures. For example, long-lived players' expected utility function $\mathbb{E}[V(X(\theta),\hat{\theta},\theta,\rho)]$ in (5.10) is not supermodular in $(\hat{\theta},\theta)$ so we do not have a supermodular game. However, our numerical procedure, to be described below, always finds a unique equilibrium matching threshold for every parameter combination.

To evaluate the equilibrium matching threshold of long-lived players, we need to compute the function V first. It is hard to find a closed-form solution to the difference equation in (5.8) and our limiting regime in Section 4 does not help. The challenge is that x can go to infinity in (5.8). Thus, numerically solving this system of difference equations exactly also appears challenging. Here, we provide a good approximation method to solve for function V heuristically.

PROPOSITION 3. *Consider $\epsilon,\rho\in[0,1]$, and fix $\theta,\hat{\theta}\in[0,\epsilon\rho]$.*

(i) There exists an upper bound B such that $V(x,\hat{\theta},\theta,\rho)\le B<\infty$ for all $x\ge 0$.

(ii) Fix $\bar{x}\in\mathbb{Z}_+$. Denote $V_{\bar{x}}(x,\hat{\theta},\theta,\rho)$ as the solution to the system of equations that solves (5.8) and (5.9) for $x<\bar{x}$ with $V_{\bar{x}}(\bar{x},\hat{\theta},\theta,\rho)=0$. We have

$$0\le V(x,\hat{\theta},\theta,\rho)-V_{\bar{x}}(x,\hat{\theta},\theta,\rho)\le \frac{B}{\left(1+\frac{\gamma}{\lambda_L}\right)^{\bar{x}-x}},\quad \forall\, 0\le x\le\bar{x}. \tag{5.12}$$

Proposition 3(ii) suggests an efficient heuristic to compute the value function V numerically, which only involves solving a system of sparse linear equations with $\bar{x}$ variables. That is, we choose a truncation point $\bar{x}$ and solve a system of equations that follow (5.8) and (5.9) for state $x\in\{0,\dots,\bar{x}-1\}$. For any $x\ge\bar{x}$, we simply set $V_{\bar{x}}(x)=0$. Furthermore, Proposition 3(ii) indicates the errors introduced by this heuristic calculation for $x\le\bar{x}$ decrease exponentially with $\bar{x}-x$. In fact, as Figure 2 suggests, the difference on the value function of using $\bar{x}=5000$ instead of $\bar{x}=500$ is negligible for all $x<490$. So we do not need to choose a very large $\bar{x}$. Moreover, according to Proposition 1, we have $X(\theta)\succeq_1 Y(0)$, where $Y(0)$ is a Poisson random variable with load $\dfrac{\lambda_L}{\gamma}$. Thus, the distribution function of $X(\theta)$ is also "light-tailed." Therefore, following Proposition 3(ii), our heuristic of setting $V_{\bar{x}}=0$ for all $x\ge\bar{x}$ for reasonably large $\bar{x}$ also has little impact on long-lived players' expected utility upon arrival, $\mathbb{E}[V(X(\theta),\hat{\theta},\theta,\rho)]$.

In the following numerical examples, we focus on function $V_{\bar{x}}$ instead of function V. That is, we compute the equilibrium matching threshold $\theta(\rho)$ as

$$\theta(\rho)\in\arg\max_{\hat{\theta}\in[0,\epsilon\rho]}\mathbb{E}[V_{\bar{x}}(X(\theta(\rho)),\hat{\theta},\theta(\rho),\rho)]. \tag{5.13}$$

We perform a fixed-point iteration algorithm to find the equilibrium matching threshold of (5.13) in our numerical procedures. Denote θ_k as the matching threshold used by players

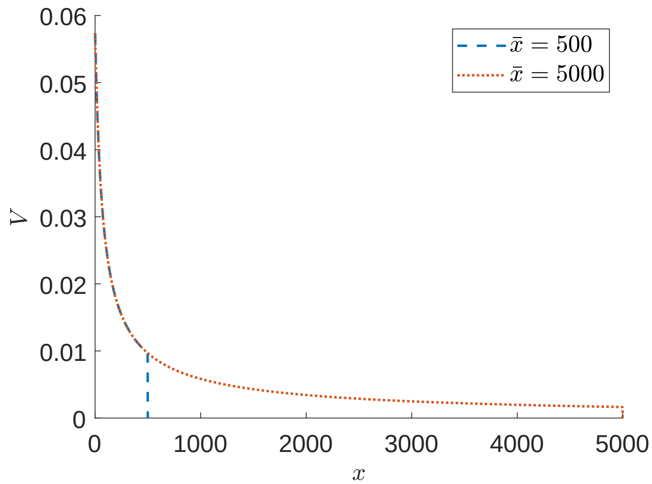

Figure 2 Function $V_{\bar{x}}$ when $\bar{x} = 500$ v.s. $\bar{x} = 5000$, with parameters $\gamma = 10$, $\epsilon = 0.1$, $\lambda_L = 20$ and $\lambda_S = 30$

other than the focal one in each iteration $k \geq 0$. Starting with $\theta_0 = \epsilon\rho$, we compute $\theta_{k+1} \in \arg\max_{\hat{\theta} \in [0, \epsilon\rho]} \mathbb{E}[V_{\bar{x}}(X(\theta_k), \hat{\theta}, \theta_k, \rho)]$ for all $k \geq 0$. In every parameter combination, we find that there is a unique maximizer in each iteration. The iterations stop and we set $\theta(\rho) = \theta_k$ when θ_k and θ_{k-1} are close enough.

As mentioned earlier in this section, if the platform sets a price ρ, long-lived players in general choose a matching threshold $\theta \leq \epsilon\rho$. Furthermore, by comparing the platform's problems in (4.1) and (5.1), if $\theta(\rho) = \epsilon\rho^*$, where ρ^* is the optimal solution to (4.1), the platform can recover the optimal social welfare rate in this decentralized system. Therefore, when long-lived players choose the matching threshold, the platform needs to inflate the price ρ to be higher than ρ^* in order to maximize social welfare.

In our numerical results, we find that the platform always sets the price at $\rho_L = 1$, the highest possible value. Given this price, long-lived players' equilibrium threshold $\theta(\rho_L)$ is still lower than the centralized optimal threshold $\epsilon\rho^*$. This observation implies that long-lived players are so picky that even if the platform uses the maximum price 1, it cannot induce the optimal centralized matching threshold. Therefore, when long-lived players design the matching threshold, the system cannot achieve optimal social welfare.

Figure 3 provides examples that illustrate the relationship between the matching thresholds and players' tolerance, departure rate, or players' arrival rates. According to Figure 3 (a), as the tolerance ϵ increases, both the centralized optimal matching threshold $\epsilon\rho^*$ and the equilibrium matching threshold $\theta(\rho_L)$ among players are increasing. Figure 3 (b) shows that as the departure rate γ (or patience level) of long-lived players increases, both matching thresholds increase as well. Figure 3 (c) demonstrates that as short-live players' arrival rate λ_S increases, both matching thresholds

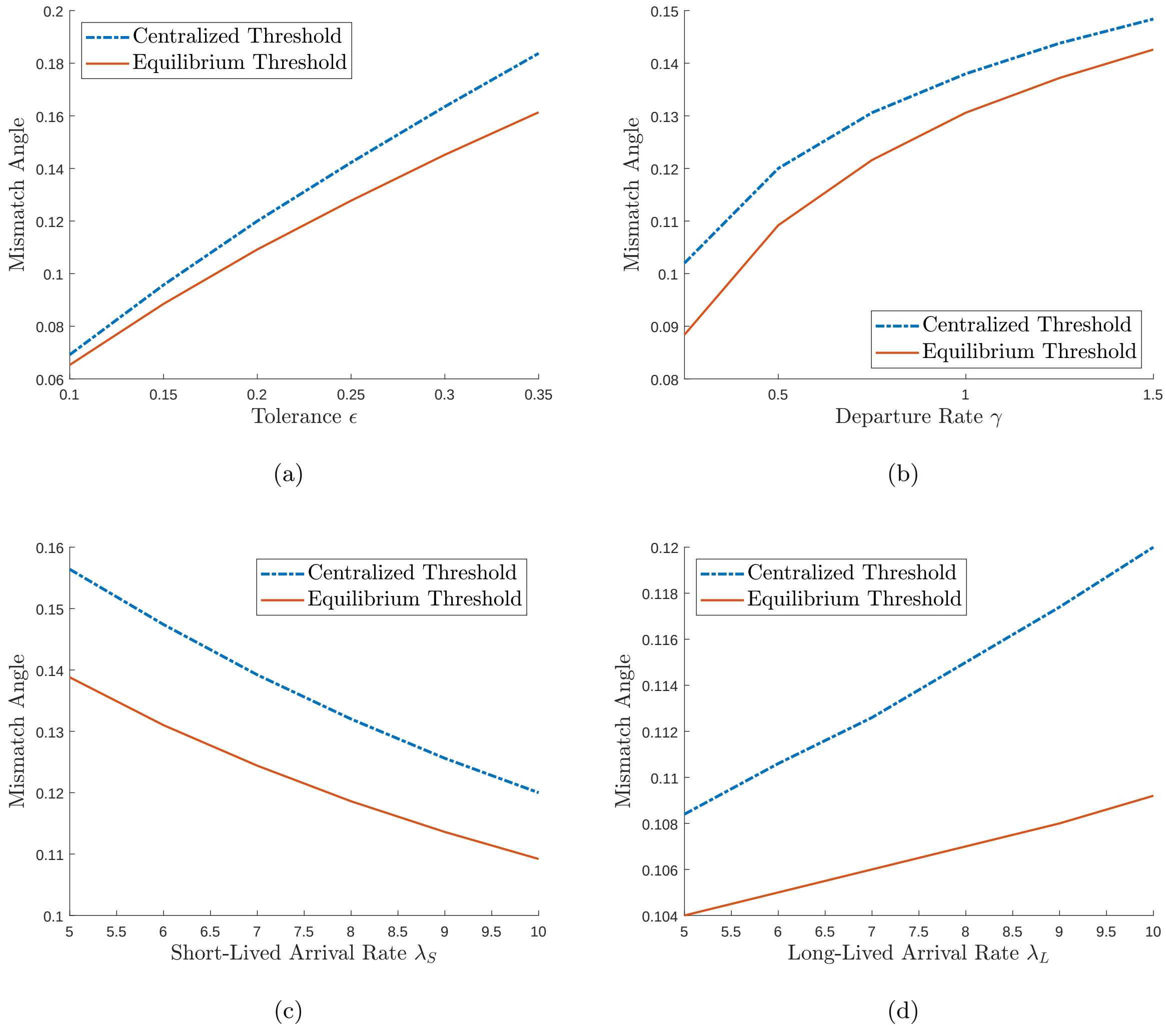

Figure 3 Centralized Matching Thresholds $\epsilon\rho^*$ v.s. Decentralized Equilibrium Matching Thresholds $\theta(\rho_L)$ with parameters a) see Table 1, b) see Table 2, c) see Table 3, and d) see Table 4

decrease. Finally, Figure 3 (d) illustrates that as long-lived players' arrival rate λ_L increases, both matching thresholds increase too. These results are fairly intuitive. The more tolerate long-lived players are (larger ϵ), the less surplus transfer from short-lived player to long-lived players are needed to induce participation, implying larger matching thresholds. The less patient long-lived players are (higher γ), the more they behave like short-lived players and choose bigger mismatch angle thresholds. In terms of arrival rates, if short-lived players arrive more frequently, long-lived players can be pickier. However, when long-lived players arrive more frequently, their competition level increases. Thus, they have to be more tolerant. Furthermore, the gap between the equilibrium matching threshold $\theta(\rho_L)$ and the centralized optimal price ρ^* appears to be increasing as player's tolerance ϵ or arrival rates λ_S and λ_L increase. The gap appears to be decreasing as players' departure rate γ (patience) increases.

Tables 1, 2, 3, and 4 (corresponding to Figure 3 (a), (b), (c), and (d), respectively) summarize the results when we change long-lived players' tolerance ϵ, their patience level (departure rate) γ, short-lived players' arrival rate λ_S and long-lived players' arrival rate λ_L, respectively. In particular, the tables report (columns from left to right):

- the optimal centralized price ρ^*;
- the change in matching rate when long-lived players choose the matching threshold comparing to centralized matching: $\Delta R := 1 - \dfrac{R(\theta(\rho_L))}{R(\epsilon\rho^*)}$;
- social welfare generated by a *successful* match under centralized matching: $sw_C = \mathbb{E}_{X(\epsilon\rho^*)}\left[\mathbb{E}_{\underline{\phi}}\left[W_{SW}(\underline{\phi}) \,|\, \underline{\phi} \le \epsilon\rho^*, X(\epsilon\rho^*)\right]\right]$;
- social welfare generated by a *successful* match when long-lived players design the matching threshold: $sw_L = \mathbb{E}_{X(\theta(\rho_L))}\left[\mathbb{E}_{\underline{\phi}}\left[W_{SW}(\underline{\phi}) \,|\, \underline{\phi} \le \theta(\rho_L), X(\theta(\rho_L))\right]\right]$;
- the change of social welfare per arrival of short-lived player (compared to centralized matching) if the platform uses centralized price when long-lived players design the matching threshold: $\Delta U(\rho^*) := \dfrac{U(\theta(\rho^*))}{U(\epsilon\rho^*)} - 1$;
- the change of social welfare per arrival of short-lived player (compared to centralized matching) if the platform uses the optimal decentralized price ρ_L when long-lived players design the matching threshold. That is, $\Delta U(\rho_L) := \dfrac{U(\theta(\rho_L))}{U(\epsilon\rho^*)} - 1$;
- the change of long-lived players' utilities if they design the matching threshold comparing to centralized matching. That is, $\Delta L := \dfrac{L(\rho_L, \theta(\rho_L))}{L(\rho^*, \epsilon\rho^*)} - 1$.

ϵ	ρ^*	$\Delta R\,(\%)$	sw_C	sw_L	$\Delta U(\rho^*)(\%)$	$\Delta U(\rho_L)$ (%)	$\Delta L(\%)$
0.1	0.692	−6.76	0.207	0.212	−3.78	−0.12	+82.3
0.15	0.638	−5.23	0.325	0.333	−4.60	−0.20	+104
0.2	0.600	−6.22	0.445	0.458	−5.06	−0.26	+120
0.25	0.569	−6.38	0.568	0.586	−5.44	−0.31	+135
0.3	0.545	−7.19	0.693	0.715	−5.64	−0.35	+148
0.35	0.525	−7.48	0.819	0.845	−5.78	−0.39	+159

Table 1 Summary for $\epsilon = 0.05$ to 0.35, $c = 3$, $\lambda_L = 10$, $\lambda_S = 10$ and $\gamma = 0.5$

γ	ρ^*	$\Delta R\,(\%)$	sw_C	sw_L	$\Delta U(\rho^*)(\%)$	$\Delta U(\rho_L)$ (%)	$\Delta L(\%)$
0.25	0.510	−10.6	0.476	0.491	−6.35	−0.45	+174
0.5	0.600	−6.81	0.445	0.458	−5.06	−0.26	+120
0.75	0.653	−5.02	0.425	0.436	−4.23	−0.17	+96.0
1	0.690	−3.73	0.409	0.419	−3.64	−0.21	+81.4
1.25	0.719	−3.07	0.396	0.405	−3.11	−0.09	+71.1
1.5	0.742	−2.53	0.384	0.391	−2.72	−0.07	+63.5

Table 2 Summary for $\gamma = 0.25$ to 1.5, $c = 3$, $\lambda_L = 10$, $\lambda_S = 10$ and $\epsilon = 0.2$

λ_S	ρ^*	$\Delta R(\%)$	sw_C	sw_L	$\Delta U(\rho^*)(\%)$	$\Delta U(\rho_L)$ (%)	$\Delta L(\%)$
5	0.782	−7.28	0.431	0.445	−2.13	−0.31	+44.6
6	0.737	−6.99	0.432	0.447	−2.85	−0.33	+59.2
7	0.696	−7.05	0.435	0.449	−3.55	−0.32	+74.6
8	0.660	−7.04	0.438	0.452	−4.16	−0.31	+89.9
9	0.628	−6.93	0.442	0.455	−4.58	−0.29	+105
10	0.600	−6.81	0.445	0.458	−5.06	−0.26	+120

Table 3 Summary for $\lambda_S = 5$ to 10, $c = 3$, $\lambda_L = 10$, $\epsilon = 0.2$ and $\gamma = 0.5$

λ_L	ρ^*	$\Delta R(\%)$	sw_C	sw_L	$\Delta U(\rho^*)(\%)$	$\Delta U(\rho_L)$ (%)	$\Delta L(\%)$
5	0.542	−2.55	0.430	0.438	−4.93	−0.06	+160
6	0.553	−3.71	0.437	0.446	−5.27	−0.09	+152
7	0.563	−4.86	0.441	0.451	−5.42	−0.13	+144
8	0.575	−5.36	0.443	0.454	−5.34	−0.16	+136
9	0.587	−5.83	0.445	0.456	−5.26	−0.20	+129
10	0.600	−6.81	0.445	0.458	−5.06	−0.26	+120

Table 4 Summary for $\lambda_L = 5$ to 10, $c = 3$, $\epsilon = 0.2$, $\lambda_S = 10$ and $\gamma = 0.5$

Although the optimal social welfare cannot be achieved under long-lived-select matching, the loss in social welfare comparing to centralized matching is minimal. The reason is that when price ρ_L is pushed to 1, the corresponding equilibrium matching threshold $\theta(\rho_L)$ is very close to $\epsilon\rho^*$, the matching threshold that achieves optimality. As we can see from Tables 1, 2, 3, and 4, when long-lived players decide the matching threshold, each successful match generates slightly higher social welfare (reported in sw_L) comparing to that of centralized matching (reported in sw_C), but the matching rate lowers by $2-10\%$ in the examples above (reported in ΔR). As a result, the social welfare generated per arrival of short-lived player is lower than that of centralized matching, but the overall losses (reported in $\Delta U(\rho_L)$) are less than 1 percent in all examples. However, if the platform does not inflate the price, e.g. simply using the optimal centralized price, there are noticeable losses in social welfare (reported in $\Delta U(\rho^*)$) comparing to using $\rho_L = 1$. Furthermore, when $\rho_L = 1$, short-lived players' utilities are also kept at 0. Thus, comparing to utilities under centralized matching processes, long-lived players' utilities (reported in ΔL) are improved significantly if they design the matching threshold.

Finally, we examine the short-lived-select matching. In this case, the platform provides each arriving short-lived player a list of long-lived player(s) who earn(s) non-negative utilities from matching with her, and she can choose the one to match with. As the matching power is given to short-lived players, the platform curbs the pickiness of long-lived players by asking for their commitment on accepting any match that brings a non-negative utility. Similar to centralized matching, the platform commits to exclude anyone who has declined a match from future matching, so it is in the best interest of long-lived players to accept purposed matches. Since short-lived players are not penalized by mismatch angles, she may pick any long-lived players who shall accept her match. In other words, given price ρ, each short-lived player effectively uses a matching threshold $\Theta(\rho) = \epsilon\rho$. If there are multiple long-lived players satisfy the threshold, the short-lived player picks the one who she has the smallest mismatch angle with[3]. As a result, the platform can simply use the optimal centralized price that maximize (4.1) in Section 4 to achieve the optimal social welfare when short-lived players chooses the matching threshold.

We summarize our findings about the decentralized matching systems. Under long-lived-select matching, long-lived players are *picky*. Namely, the platform's desired matching threshold cannot be enforced, because long-lived players choose an equilibrium matching threshold smaller than the one induced by price $\epsilon\rho$. To counter long-lived players' behavior, the platform needs to inflate its price purposely to recover social welfare. However, this price inflation hurts short-lived players by reducing their surplus in each match and also dramatically increases long-lived players' surplus as we can see from Tables 1, 2, 3, and 4. Although social welfare in this setting can be close to the optimal centralized social welfare, the platform should exercise caution with this matching system due to its extreme welfare allocation between short- and long-lived players. In addition, letting-long-lived players decide the matching threshold also reduces the overall matching rate by $2 - 10\%$ in our numerical study. In comparison, with short-lived-select matching, the platform can simply use the optimal centralized price to achieve the optimal social welfare. The corresponding matching rate and welfare allocation in this case are also the same as those under the optimal centralized matching.

6. Extension: Heterogeneous Players

So far in this paper, we assume that short-lived players' benefit upon matching, u, is fixed and observable to the platform. Now we consider short-lived players' benefit per match is heterogeneous and not observable to the platform. That is, u is a random variable, which follows a known distribution. For example, consider a carpooling scenario where short-lived players are riders who have many outside options ranging from private cars to public transportations. The difference in their outside options imply heterogeneous values of u. For simplicity we focus on heterogeneity on short-lived players' side only.

In this section, we assume that random variable u has mean $\bar{u}$, and thus we modify the definition of ϵ and ρ representing players' average tolerance and normalized price to:

$$\bar{\epsilon} := \frac{\bar{u}}{c}, \text{ and } \bar{\rho} := \frac{P}{\bar{u}}, \text{ respectively.} \tag{6.1}$$

Furthermore, we assume

$$u = \bar{u} + sc\bar{\epsilon}, \tag{6.2}$$

where s is a uniform random variable on $[-\sigma, \sigma]$, with $0 < \sigma \leq 1$. Therefore, u is a uniform random variable on $[\bar{u}(1-\sigma), \bar{u}(1+\sigma)]$ and σ characterizes its spread around the mean. As assumed in Sections 4 and 5, the platform controls the price $\bar{\rho}$. The platform or players can act as decision makers who choose the matching threshold $\Theta(\bar{\rho})$ on mismatch angles, similar to previous sections.

Under each price $\bar{\rho}$ set by the platform, a short-lived player's utility is

$$u - P = u - \bar{u}\bar{\rho} = u - \bar{\rho}c\bar{\epsilon} = c\bar{\epsilon}(1 + s - \bar{\rho}), \tag{6.3}$$

which follows (6.1) and (6.2). Thus, a short-lived player participates in matching if and only if $u \geq P$, or, equivalently, $s \geq \bar{\rho} - 1$. Given $\sigma \in [0,1]$, the probability that an arriving short-lived player shall participate in the matching process is

$$\mathbb{P}(s \geq \bar{\rho} - 1) = \begin{cases} 0, & \text{if } \bar{\rho} > 1 + \sigma, \\ \dfrac{\sigma + 1 - \bar{\rho}}{2\sigma}, & \text{if } 1 - \sigma \leq \bar{\rho} \leq 1 + \sigma, \\ 1, & \text{if } 0 \leq \bar{\rho} < 1 - \sigma. \end{cases} \tag{6.4}$$

As a result, when there are x long-lived players on the platform and the matching threshold is θ, the matching probability in this new system is $\mathbb{P}(s \geq \bar{\rho} - 1)p_x(\theta)$, where $p_x(\theta)$ is defined in (3.5).

6.1. Centralized matching.

Note, without heterogeneity, the benefit generated in each match is exactly u, which is independent of the price, so the platform can only affect the social welfare by its pricing that induces a threshold on the mismatch angle. However, with heterogeneity of u, there are two ways the platform can affect the social welfare. First, the centralized pricing decision, denoted as $\bar{\rho}_C$, may introduce an entry barrier on short-lived players, which affects the benefit generated per match. Second, a centralized matching threshold, denoted as $\bar{\theta}_C$, on mismatch angle affects matching rate and quality. Therefore, unlike Section 4, the platform may set a matching threshold other than $\bar{\theta}_C = \bar{\epsilon}\bar{\rho}_C$ to maximize social welfare. Similar to previous sections, we refer to the expected social welfare generated per underlying arrival of short-lived players (with rate λ_S) as the *social welfare*. That is, the platform's centralized problem now becomes:

$$\max_{\bar{\rho}_C \in [0, 1+\sigma], \bar{\theta}_C \in [0, \bar{\epsilon}\bar{\rho}_C]} \mathbb{E}_{X(\bar{\rho}_C, \bar{\theta}_C), s}[\bar{h}_{\bar{\epsilon}}(X(\bar{\rho}_C, \bar{\theta}_C), \bar{\theta}_C, s)\mathbb{I}\{s \geq \bar{\rho} - 1\}], \tag{6.5}$$

where function $\bar{h}_\epsilon$ is defined, similarly to function h_ϵ in (4.2), as

$$\bar{h}_{\bar{\epsilon}}(x, \theta, s) = \mathbb{E}_{\underline{\phi}}\left[(u - c\underline{\phi})\mathbb{I}\{\underline{\phi} \leq \theta\} \,|\, x\right] = \mathbb{E}_{\underline{\phi}}\left[c(\bar{\epsilon}(1+s) - \underline{\phi})\mathbb{I}\{\underline{\phi} \leq \theta\} \,|\, x\right], \tag{6.6}$$

and, with slight abuse of notation, random variable $X(\bar{\rho}_C, \bar{\theta}_C)$ follows the steady state distribution of a Birth-Death process with birth rate λ_L and death rate $\lambda_S \mathbb{P}(s \geq \bar{\rho}_C - 1)p_x(\bar{\theta}_C)$. The constraint $\bar{\theta}_C \leq \bar{\epsilon}\bar{\rho}_C$ in (6.5) ensures players accept each match initiated by the platform.

Similar to Section 4, we introduce a myopic policy that maximizes the matching rate while generating non-negative social welfare per match. Unlike the case without heterogeneity, the matching rate needs to incorporate the entry barrier. That is, we define the expected matching rate under price ρ and matching threshold θ as:

$$\lambda_S \mathbb{P}(s \geq \rho - 1)\mathbb{E}_{X(\rho,\theta)}[p_{X(\rho,\theta)}(\theta)] = \lambda_S \mathbb{P}(s \geq \rho - 1)\mathbb{E}_{X(\rho,\theta)}[1 - (1-\theta)^{X(\rho,\theta)}], \tag{6.7}$$

where function $p_x(\cdot)$ is defined in (3.5) representing the matching probability. From (6.7), we immediately observe that the matching rate is increasing with respect to θ, for any ρ. Recall that $\theta \leq \bar{\epsilon}\rho$, because a long-lived player's' utility may be negative in a match, otherwise. Thus, in a myopic matching policy, we can simply set the matching threshold $\theta = \bar{\epsilon}\rho$ and find the price ρ that maximize (6.7). Therefore, we define the myopic price $\bar{\rho}_M$ such that

$$\bar{\rho}_M \in \underset{\rho \in [0,1+\sigma]}{\arg\max}\, \lambda_S \mathbb{P}(s \geq \rho - 1)\mathbb{E}_{X(\rho,\bar{\epsilon}\rho)}[p_{X(\rho,\theta)}(\theta)]. \tag{6.8}$$

In our numerical studies, the myopic matching policy with price $\bar{\rho}_M$ and threshold $\bar{\epsilon}\bar{\rho}_M$ is still near optimal when $\bar{\epsilon}$ is small, just like the case without heterogeneity in Section 4. However, the myopic matching policy may not induce extremely unbalanced welfare allocation among players. The reason is that with heterogeneity, increasing the price charged to short-lived players shall increase the entry barrier, which decreases the matching rate. Therefore, the myopic price may be even smaller than the optimal one. As a result, although some short-lived players are screened off by the entry barrier, many others with high benefit can earn great surplus.

6.2. Decentralized matching.

In decentralized matching systems, we first consider long-lived-select matching. With players' heterogeneity, we can modify the platform's problem in (5.1) to

$$\max_{\bar{\rho} \in [0,1+\sigma]} \mathbb{E}_{X(\bar{\theta}_L(\bar{\rho})),s}[\bar{h}_{\bar{\epsilon}}(X(\bar{\theta}_L(\bar{\rho})), \bar{\theta}_L(\bar{\rho}), s)\mathbb{I}\{s \geq \bar{\rho} - 1\}], \tag{6.9}$$

where function $\bar{h}_\epsilon$ is defined in (6.6) and $\bar{\theta}_L(\bar{\rho})$ is long-lived players' equilibrium matching threshold. When there is no confusion, we write $\bar{\theta}_L(\bar{\rho})$ as $\bar{\theta}_L$. The platform's price not only screens off short-lived players with low benefit, but also induces a game between long-lived players just like in Section 5. Thus, we leverage on results of Proposition 3 and extend our numerical procedure. We present the procedure on finding the equilibrium matching threshold $\bar{\theta}_L$ in the Appendix B.4, which is based on the procedure introduced in Section 5. Before moving on, note that the threshold designed by short-lived players is induced by the platform's price. Thus, although there is heterogeneity among them, their myopic decision leads to the same threshold under the same price. We leave the case where there is heterogeneity among long-lived players (e.g. they have different mismatch cost) as a future research direction since thresholds decided by them depend on their heterogeneous values.

Under short-lived-select matching, short-lived players simply set the threshold as $\Theta(\bar{\rho}) = \bar{\epsilon}\bar{\rho}$. The reason is that, similar to Section 5, each short-lived player is myopic and chooses the closest long-lived player who accepts the mismatch cost, i.e., their mismatch angle is no greater than $\bar{\epsilon}\bar{\rho}$. Therefore, the platform's problem of maximizing social welfare is

$$\max_{\bar{\rho}\in[0,1+\sigma]} \mathbb{E}_{X(\bar{\epsilon}\bar{\rho}),s}[\bar{h}_{\bar{\epsilon}}(X(\bar{\epsilon}\bar{\rho}),\bar{\epsilon}\bar{\rho},s)\mathbb{I}\{s \geq \bar{\rho}-1\}], \tag{6.10}$$

where function $\bar{h}_\epsilon$ is defined in (6.6) and the matching threshold is $\bar{\epsilon}\bar{\rho}$. Here, the platform only chooses the price $\bar{\rho}$.

In order to better illustrate our results, we present a numerical study using an example on ride-sharing similar to a context of DiDi Hitch. Riders and drivers depart from the center of a circle with a radius of 24km (15miles)[4] and have their destinations uniformly distributed on the edge of the circle. In the numerical study, we assume the pricing rule of taxi service represents the mean of riders' benefit, $\bar{u}$, since it is a common medium-priced alternative for many people. Take daytime taxi service in Beijing for example, each rider needs to pay a fare equivalent to 1.4 dollar for the first three kilometers, and 0.28 dollar per kilometer afterwards. That is, a 24km trip leads to an average benefit of $\bar{u} = 1.4 + 0.28 \times 21 = 7.28$ dollars. In other words, a rider saves the taxi fare $\bar{u}$ if successfully matched in DiDi Hitch. On the drivers' side, we assume they are daily commuters with an outside opportunity cost of 15 dollars per hour. The travel cost is estimated as the fuel cost. Consider a car consuming 10 liter per 100 kilometers and the gas price is \$1.02 per liter. Using rush hour average travel speed of 36km/hr (22.5mph), we consider drivers' costs with the unit of dollar per kilometer, as reported in Table 5. Since the detour is the length of the arc between a driver and a rider's destinations on a circle with 24km radius, we have the average tolerance of a driver as,

$$\bar{\epsilon} = \frac{\bar{u}}{c} = \frac{\$7.28}{\$(0.417+0.102)/km \times \pi \times 24km} \approx 0.186. \tag{6.11}$$

Driver Characteristics	
Travel speed	36km/hr
Opportunity cost	\$0.417/km
Travel cost	\$0.102/km
Total cost	\$0.519/km

Table 5 Driver Characteristics

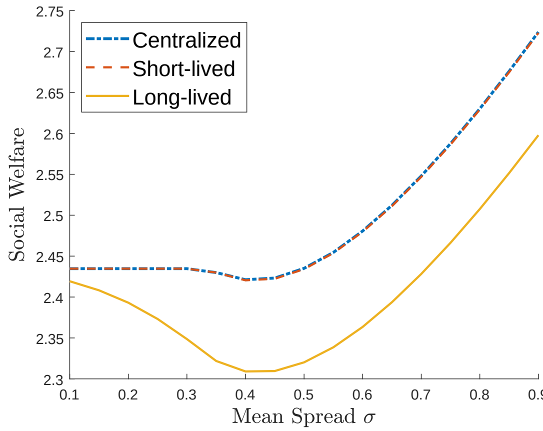

(a)

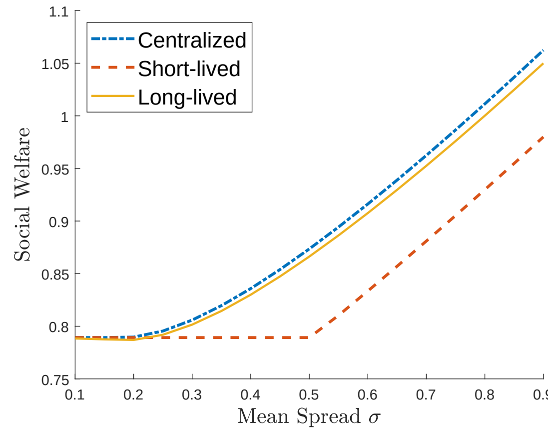

(b)

Figure 4 Social welfare comparison, $\lambda_L = 5$, $\bar{\epsilon} = 0.186$, $c = 39.13$, $\gamma = 0.5$, (a) $\lambda_S = 5$, (b) $\lambda_S = 30$

Figure 4 shows the social welfare (in dollars) generated per *underlying* arrival (with rate λ_S, without entry barrier) of short-lived player with respect to different mean spread σ when the platform, short- and long-lived players decide the matching threshold, respectively. In Figure 4(a), with a relatively low λ_S, each arrival of short-lived players generates between 2.44 and 2.80 dollars of social welfare in either centralized or short-lived-select matching. Under long-lived-select matching, each arrival generates between 2.31 to 2.65 dollars. In Figure 4(b), with a relatively large λ_S, each arrival of short-lived players generates between 0.78 to 1.07 dollars of social welfare in centralized matching. Under short- and long-lived-select matching, social welfare can be up to 1.05, and 0.97 dollars per arrival, respectively.

Furthermore, from Figure 4(a), we observe that social welfare is not monotone with respect to σ when λ_S is relatively small. In particular, if the platform or short-lived players decide the matching threshold, social welfare stays the same for small values of σ. As σ increases, social welfare first decreases then increases for large values of σ. If long-lived players decide the matching threshold, we observe similar trends except there is no region of σ that social welfare stays the same. Note the difference in social welfare under centralized or short-lived-select matching is negligible. However, there is a noticeable gap in social welfare under long-lived-select matching, especially when σ is large. From Figure 4(b), we observe that social welfare is mostly increasing with respect to σ when λ_S is relatively large. Finally, the gap in social welfare between centralized and short-lived-select matching is significant, especially when σ is large, whereas long-live-select matching almost achieves the optimal centralized social welfare in this case.

Tables 6 and 7 complement Figure 4 by reporting (columns from left to right):

- the optimal centralized price $\bar{\rho}_C$;

σ	$\bar\rho_C$	$\bar\rho_S$	$\bar\rho_L$	$\bar\theta_C$	$\bar\theta_S$	$\bar\theta_L$	sw_C	sw_S	sw_L	SW_S (%)	SW_L (%)
0.1	0.700	0.700	0.900	0.130	0.130	0.114	4.933	4.933	5.216	0	−0.63
0.2	0.700	0.700	0.800	0.130	0.130	0.104	4.933	4.933	5.387	0	−1.71
0.3	0.700	0.700	0.700	0.130	0.130	0.094	4.933	4.933	5.570	0	−3.53
0.4	0.646	0.646	0.697	0.120	0.120	0.096	5.286	5.286	5.894	0	−4.63
0.5	0.681	0.681	0.733	0.127	0.127	0.103	5.697	5.697	6.285	0	−4.72
0.6	0.716	0.716	0.773	0.133	0.133	0.110	6.104	6.104	6.692	0	−4.72
0.7	0.752	0.752	0.810	0.140	0.140	0.117	6.510	6.510	7.100	0	−4.70
0.8	0.787	0.787	0.840	0.146	0.146	0.122	6.916	6.916	7.502	0	−4.66
0.9	0.822	0.822	0.884	0.153	0.153	0.129	7.322	7.322	7.928	0	−4.63

Table 6 Summary for $\sigma = 0.1$ to 0.9, $\lambda_S = \lambda_L = 5$, $\gamma = 0.5$, $c = 39.13$ and $\bar\epsilon = 0.186$

σ	$\bar\rho_C$	$\bar\rho_S$	$\bar\rho_L$	$\bar\theta_C$	$\bar\theta_S$	$\bar\theta_L$	sw_C	sw_S	sw_L	SW_S (%)	SW_L (%)
0.1	0.356	0.356	0.900	0.066	0.066	0.062	5.246	5.246	5.412	0	−0.11
0.2	0.822	0.356	0.829	0.068	0.066	0.061	5.318	5.246	5.627	−0.06	−0.35
0.3	0.859	0.356	0.869	0.084	0.066	0.069	5.718	5.246	6.101	−2.07	−0.54
0.4	0.905	0.356	0.916	0.093	0.066	0.076	6.112	5.246	6.559	−5.57	−0.70
0.5	0.952	0.356	0.969	0.101	0.066	0.082	6.510	5.246	7.020	−9.64	−0.83
0.6	1.003	0.863	1.102	0.109	0.160	0.088	6.905	4.639	7.474	−9.03	−0.94
0.7	1.057	0.919	1.107	0.117	0.171	0.093	7.310	5.010	7.931	−8.46	−1.03
0.8	1.111	0.975	1.113	0.122	0.181	0.099	7.718	5.363	8.392	−8.06	−1.11
0.9	1.168	1.032	1.192	0.128	0.192	0.104	8.123	5.703	8.867	−7.76	−1.17

Table 7 Summary for $\sigma = 0.1$ to 0.9, $\lambda_S = 30$, $\lambda_L = 5$, $\gamma = 0.5$, $c = 39.13$ and $\bar\epsilon = 0.186$

- the optimal decentralized price $\bar\rho_S$ when short-lived players decide the matching threshold;
- the optimal decentralized price $\bar\rho_L$ when long-lived players decide the matching threshold;
- the optimal centralized matching threshold $\bar\theta_C$;
- the decentralized matching threshold $\bar\theta_S := \bar\epsilon\bar\rho_S$ decided by short-lived players;
- the equilibrium matching threshold $\bar\theta_L$ decided by long-lived players;
- social welfare generated by a *successful* match under centralized matching:

$$sw_C := \mathbb{E}_{X(\bar\theta_C),s}\left[\mathbb{E}_{\underline{\phi}}\left[c(\bar\epsilon(1+s)-\underline{\phi})\,|\,\underline{\phi}\le\bar\theta_C, X(\bar\theta_C)\right]\,\middle|\, s\ge\bar\rho_C-1\right];$$

- social welfare generated by a *successful* match when short-lived players decide the matching threshold:

$$sw_S := \mathbb{E}_{X(\bar\theta_S),s}\left[\mathbb{E}_{\underline{\phi}}\left[c(\bar\epsilon(1+s)-\underline{\phi})\,|\,\underline{\phi}\le\bar\theta_S, X(\bar\theta_S)\right]\,\middle|\, s\ge\bar\rho_S-1\right];$$

- social welfare generated by a *successful* match when long-lived players decide the matching threshold:

$$sw_L := \mathbb{E}_{X(\bar\theta_L),s}\left[\mathbb{E}_{\underline{\phi}}\left[c(\bar\epsilon(1+s)-\underline{\phi})\,|\,\underline{\phi}\le\bar\theta_L, X(\bar\theta_L)\right]\,\middle|\, s\ge\bar\rho_L-1\right];$$

- the change in social welfare per underlying arrival of short-lived player (compared to centralized matching) if short-lived players decide the matching threshold:

$$SW_S := \frac{\mathbb{E}_{X(\bar\epsilon\bar\rho_S),s}\bar h_{\bar\epsilon}(X(\bar\epsilon\bar\rho_S),\bar\epsilon\bar\rho_S,s)\mathbb{I}\{s\ge\bar\rho_S-1\}}{\mathbb{E}_{X(\bar\theta_C),s}\bar h_{\bar\epsilon}(X(\bar\theta_C),\bar\theta_C,s)\mathbb{I}\{s\ge\bar\rho_C-1\}}-1;$$

- the change in social welfare per underlying arrival of short-lived player (compared to centralized matching) if long-lived players decide the matching threshold:

$$SW_L := \frac{\mathbb{E}_{X(\bar{\theta}_L),s}\bar{h}_{\bar{\epsilon}}(X(\bar{\theta}_L),\bar{\theta}_L,s)\mathbb{I}\{s \geq \bar{\rho}_L - 1\}}{\mathbb{E}_{X(\bar{\theta}_C),s}\bar{h}_{\bar{\epsilon}}(X(\bar{\theta}_C),\bar{\theta}_C,s)\mathbb{I}\{s \geq \bar{\rho}_C - 1\}} - 1.$$

Based on observations in Figure 4, the underlying arrival rate of short-lived players λ_S has significant impacts on social welfare and the decision of who should decide the matching threshold. Note that, fixing λ_L and γ, λ_S can be interpreted as the market thickness of short-lived players since they never stay in the matching system. Thus, we use Tables 6 and 7 to further reveal the connection between our observations and the market thickness.

First, we focus on Table 6 where short-lived players' side of the market is relatively thin ($\lambda_S = 5$). Note that in this case, when σ is small, the optimal centralized price $\bar{\rho}_C$ and short-lived price $\bar{\rho}_S$ for any σ are no greater than $1-\sigma$. For example, in Table 6, when $\sigma \leq 0.3$, we have $\bar{\rho}_C = \bar{\rho}_S = 0.700 < 1-\sigma$. According to (6.4), there is no entry barrier on short-lived players at all. As the prices remain the same and there is no screening of short-lived players, the social welfare with different σ remains the same, which is consistent with our observation in Figure 4(a). However, if long-lived players decide the matching threshold, we observe that both the price $\bar{\rho}_L$ and the equilibrium threshold $\bar{\theta}_L$ are decreasing, as σ increases on $[0, 0.3]$. The reason is that long-lived players are picky so the platform has to inflate the price to be equal to the largest price that does not induce an entry barrier on short-lived players, which decreases as σ increases.

When σ is medium, e.g. $\sigma = 0.4$, the optimal prices and matching thresholds are smaller than those for $\sigma \leq 0.3$, under centralized or short-lived-select matching. The reason is that if the platform still sets a high price, there will be an entry barrier for short-lived players. Using a smaller price leads to a lower (or even no) entry barrier that allows more short-lived players to participate in the matching process, i.e., creating a thicker market. However, a lower price also induces a smaller matching threshold, which decreases the matching rate and may hurt the overall social welfare. This also explains the decrease of social welfare in Figure 4(a) when σ is medium. When σ is large (e.g. $\sigma > 0.4$), the optimal prices increases. This is because the platform finds it more beneficial to sacrifice some short-lived players with low benefits to ensure more players with high benefits are matched. Therefore, as we can see in Figure 4(a), social welfare is increasing in σ when σ is large, no matter who decides the matching threshold.

Moreover, when short-lived players' side of the market is relatively thin, we have $\bar{\rho}_C = \bar{\rho}_S$ and $\bar{\theta}_C = \bar{\theta}_S = \bar{\epsilon}\bar{\rho}_C$ for all values of σ. Thus, centralized matching is equivalent to letting short-lived players decide the matching threshold, just like in Section 5. However, when long-lived players decide the matching threshold, we observe noticeable loss in social welfare. The reason is that long-lived players are too picky and price inflation is not as effective as when there is no heterogeneity,

$\sigma = 0$. Although the social welfare generated per successful matching (sw_L) is slightly higher when long-lived players decide the matching threshold comparing to that of the other two systems, the matching rate is lower since $\bar{\theta}_L < \bar{\theta}_C = \bar{\theta}_S$ so the overall social welfare is lower. For example, when $\sigma \geq 0.35$ in Table 6, we have gaps close to 5 percent (in SW_L) while we never observe a gap greater than 1 percent in Section 5. With heterogeneity, price inflation may induce an entry barrier on short-lived players. When λ_S is small, such an entry barrier may greatly weaken the thickness of short-lived players, resulting in lower social welfare.

Next, we turn our attention to Table 7 where short-lived players' side of the market is relatively thick ($\lambda_S = 30$). In this case, creating an entry barrier is less of a concern. The reason is that an entry barrier allows short-lived players with high benefits joining the platform but the platform still preserves a decent thickness of short-lived players due to their high underlying arrival rate. Thus, the platform can use higher prices when σ increases, comparing to the case where short-lived players' side of the market is relatively thin. Take centralized prices for example: $\bar{\rho}_C$ does not increase until $\sigma > 0.4$ when $\lambda_S = 5$. However, when $\lambda_S = 30$, $\bar{\rho}_C$ is non-decreasing for all σ. Furthermore, the platform and long-lived players all use relatively small thresholds, which not only lowers the mismatch penalty but also decreases the matching rate. This is because the matching quality is more important than the matching rate when short-lived players arrive frequently. However, short-lived players' myopic behavior leads to much larger matching thresholds $\bar{\theta}_S$ comparing to either $\bar{\theta}_C$ or $\bar{\theta}_L$ under high prices, and this leads to much lower matching quality. For example, when $\sigma = 0.6$, each successful match generates $sw_S = 4.639$, $sw_C = 6.905$, and $sw_L = 7.474$ dollars under short-lived-select, centralized, and long-lived-select matching, respectively. This means a gap of almost 10 percent in social welfare (reported in SW_S) under short-lived-select matching comparing to centralized matching. In contrast, letting long-lived players choose the matching threshold only induces a gap around 1 percent. Since long-lived players are very picky, they use much smaller matching thresholds comparing to that of short-lived players when $\sigma \geq 0.5$. In comparison, short-lived players' myopic behavior jeopardizes the matching quality, which reduces social welfare.

In Figure 5, we plot the optimal social welfare per underlying arrival of short-lived player against λ_S when short- and long-lived players choose the matching threshold, respectively. Both curves decrease in λ_S because as short-lived players arrive more frequently, higher entry barriers are placed on them, which decreases the expected social welfare generated per underlying arrival of short-lived player. Note that there is a threshold in the short-lived players' arrival rate λ_S such that it maximizes the social welfare by letting short-lived players decide the matching threshold if and only if λ_S is less than the threshold.

To summarize our findings, we recommend the platform use short-lived-select matching if either the market of short-lived players is relatively thin or the level of heterogeneity is low, i.e. λ_S or σ

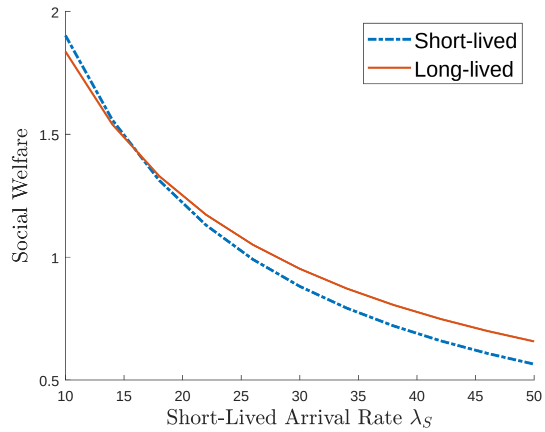

Figure 5 Social Welfare Comparison with parameters $\lambda_L = 5$, $\gamma = 0.5$, $c = 39.13$, $\bar{\epsilon} = 0.186$ and $\sigma = 0.7$

is small, respectively. In this case, the decentralized system can still achieve optimal centralized social welfare just like in Section 5 without heterogeneity. When short-lived players arrive more frequently and their heterogeneity is large, i.e., both λ_S and σ are large, using long-lived-select matching is the better choice in decentralized systems. In this case, long-lived players being picky is beneficial, because it induces a smaller matching threshold comparing to the short-lived-select matching threshold, similar to the centralized matching threshold.

7. Conclusion

We build a novel model of an online marketplace (e.g., carpooling and gig-job hunting), where one side of short-lived players are matched with the other side of long-lived players, and matching quality is an important issue. We apply this model to study pricing and matching mechanisms at Didi Hitch, the largest commuter carpooling platform in China, to address some design questions that concern the platform operator.

Our research shows how pricing and matching mechanisms affect the tradeoff between matching rate and quality, and affect the welfare allocation across the two sides of players. In particular, for the scenario with homogeneous outside option of riders (e.g., routes with taxi as the dominating alternative transportation mode), we show that the platform can have a range of nearly optimal prices that give varying weights on matching rate and matching quality and produce varying welfare allocation scenarios across the two sides of players (i.e., drivers and riders). The specific choice of matching mechanism in this spectrum has small impact on total social welfare but has significant impact on the rate-versus-quality balance and welfare allocation. For example, the platform's status-quo long-lived-select mechanism, favors matching quality over matching rate and reduces short-lived players' (riders') welfare; the platform's alternative mechanism, which restricts drivers' selection, would perform close to the centralized optimal mechanism and do a better job to balance two sides' welfare allocation.

In contrast, for the scenario with heterogeneous outside options of riders (e.g., routes with multiple comparable transportation modes), the platform needs to optimize the price given the chosen matching mechanism, and the chosen matching mechanism can significantly affect total social welfare. We recommend to use the alternative mechanism (that has riders to select drivers) if either the market of riders is relatively thin or the level of heterogeneity is low but, otherwise, to keep the status-quo mechanism (that has drivers to select riders).

Therefore, our research contributes to the related literature on two-sided matching by first comparing two decentralized matching mechanisms and addressing the important design question which side should be let to select the other? Our research also highlights the role of pricing in different scenarios. In the homogeneous outside option scenario, pricing plays the critical role of balancing matching quantity and quality, and also affecting welfare allocation; but it does not significantly affect total welfare in the sense that there exists a range of nearly optimal prices. However, in the heterogeneous outside option scenario, pricing has the additional marginal effect to deter short-lived players with relatively low valuations, and thus can significantly affect the total welfare.

Our model and analysis potentially provide a useful technical framework for related problems in online matching. For example, we provide a simple closed-form approximation on the centralized optimal matching threshold, which is based on a $M/M/\infty$ model of the long-lived players in the market. For the game theory analysis of long-lived players' strategy in the decentralized matching mechanism, we provide an efficient and accurate numerical procedure to compute the equilibrium.

Although this paper focuses on welfare maximization, our analytical approaches developed in Sections 4 and 5 can readily extend to focus on the objective of profit maximization. Our model can also easily extend to allow the two sides' players to split mismatch costs, for example, to allow a hirer and a gig-worker to negotiate the work order and thus share the mismatch cost. In certain centralized matching settings, the platform may be able to use additional information, such as the current number of players on the market, as well as their positions, to dynamically decide matching criteria. We leave this complexity out of this paper, partly because the platform in our motivation found it too costly and complex to develop and implement; however, it remains interesting to investigate if and when dynamic matching control significantly improves beyond the simple threshold policy that we considered in Section 4.

Endnotes

1. We show that the stationary distribution always exists in Appendix B.1.

2. For any two random variables X and Y with cumulative probability distributions F_X and F_Y respectively, X first-order stochastically dominates X, written $X \succeq_1 Y$, if and only if $F_X(k) \leq F_Y(k)$ for all k.

3. This tie-breaking rule has many applications in reality. For example, a rider on DiDi Hitch chooses the driver with whom she has the shortest distance between destinations, since she may want to minimize carbon emission.

4. We use the average driving distance in DiDi Hitch (see, e.g., Liu et al. 2019).

References

Adan, I. and Weiss, G. (2012). Exact FCFS matching rates for two infinite multi-type sequences. *Operations Research*, 60(2):475–489.

Afèche, P., Diamant, A., and Milner, J. (2014). Double-sided batch queues with abandonment: Modeling crossing networks. *Operations Research*, 62(5):1179–1201.

Akan, M., Alagoz, O., Ata, B., and Erenay, F. (2012). A broader view of designing the liver allocation system. *Operations Research*, 60(4):757–770.

Akbarpour, M., Li, S., and Gharan, S. (2019). Thickness and information in dynamic matching markets. *Working Paper.*

Allon, G., Bassamboo, A., and Çil, E. (2012). Large-scale service marketplaces: The role of the moderating firm. *Management Science*, 58(10):1854–1872.

Anderson, R., Ashlagi, I., Y., K., and Gamarnik, D. (2017). Efficient dynamic barter exchange. *Operations Research*, 65(6):1446–1459.

Arnosti, N., Johari, R., and Kanoria, Y. (2020). Managing congestion in matching markets. *Manufacturing & Service Operations Management, forthcoming.*

Ashlagi, I., Burq, M., Jaillet, P., and Manshadi, V. (2019). On matching and thickness in heterogeneous dynamic markets. *Operations Research*, 67(4):927–949.

Büke, B. and Chen, H. (2017). Fluid and diffusion approximations of probabilistic matching system. *Queueing System*, 86:1–33.

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

Chen, Y. and Hu, M. (2020). Pricing and matching with forward-looking buyers and sellers. *Manufacturing & Service Operations Management*, 22(4):717–734.

Feng, G., Kong, G., and Wang, Z. (2020). We are on the way: Analysis of on-demand ride-hailing. *Manufacturing & Service Operations Management, forthcoming.*

Gurvich, I. and Ward, A. R. (2014). On the dynamic control of matching queues. *Stochastic Systems*, 2(4):479–523.

Iglehart, D. (1965). Limiting diffusion approximations for the many server queue and repairement problem. *Journal of Applied Probability*, 2(2):429–411.

Kanoria, Y. and Saban, D. (2020). Facilitating the search for partners on matching platforms. *Management Science, forthcoming.*

Liu, X., Yang, C., and Wan, Z. (2019). The efficiency of a dynamic decentralized two-sided matching market. *Working Paper.*

Özkan, E. and Ward, A. (2019). Dynamic matching for real-time ridesharing. *Stochastic Systems, forthcoming.*

Pavan, A. and Gomes, R. (2019). Price customization and targeting in matching markets. *Working paper.*

Salop, S. C. (1979). Monopolistic competition with outside goods. *The Bell Journal of Economics*, 10(1):141–156.

Su, X. and Zenios, S. (2006). Recipient choice can address the efficiency-equity and trade-off in kidney transplantation: A mechanism design model. *Management Science*, 52(11):1647–1660.

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

Yang, P., Iyer, K., and Frazier, P. (2016). Mean field equilibria for competitive exploration in resource sharing settings. *Proc. 25th Internat. World Wide Web Conf.*, pages 177–187.

Zenios, S., Chertow, G., and Wein, L. (2000). Dynamic allocation of kidneys to candidates on the transplant waiting list. *Operations Research*, 48(4):549–569.

Appendix

A. Proofs for Major Results

Proof of Lemma 1. We prove this result by showing that $\psi(x) := \epsilon\rho x - [1-(1-\epsilon\rho)^x]$ is non-negative for all $x \geq 1$.

The first and second order derivatives of $\psi(x)$ are

$$\frac{d\,\psi(x)}{d\,x} = \epsilon\rho + (1-\epsilon\rho)^x \ln(1-\epsilon\rho), \text{ and } \frac{d^2\,\psi(x)}{d\,x^2} = (1-\epsilon\rho)^x(\ln(1-\epsilon\rho))^2, \text{ respectively.}$$

Since $0 \leq \epsilon\rho \leq 1$, we have $\frac{d^2\,\psi(x)}{d\,x^2} \geq 0$, and, therefore,

$$\frac{d\,g_1(x)}{d\,x} \geq \left.\frac{d\,g_1(x)}{d\,x}\right|_{x=1} = \xi(\rho), \text{ where } \xi(\rho) := \epsilon\rho + (1-\epsilon\rho)\ln(1-\epsilon\rho).$$

Function $\xi(\cdot)$ is decreasing because

$$\frac{d\,\xi(\rho)}{d\,\rho} = -\epsilon\ln(1-\epsilon\rho) \geq 0,$$

for $\epsilon\rho \in [0,1]$. Therefore, we have

$$\frac{d\,\psi(x)}{d\,x} \geq \xi(0) = 0.$$

As the result, we have $\psi(\cdot)$ is an increasing function. Since $\psi(1)=0$, we reach the final result that $\psi(x) \geq 0$ for all $x \geq 1$.

□

We show a more general result in order to prove Proposition 1.

LEMMA 3. *Consider two Birth-Death Processes $X_1 = \{X_1(t), t \geq 0\}$ and $X_2 = \{X_2(t), t \geq 0\}$ with the same arrival rate λ. Process X_1 and X_2 have departure rates $\mu_1(x)$ and $\mu_2(x)$ such that $\mu_1(x) \geq \mu_2(x)$ for all states $x \geq 1$. Denote $\bar{X}_1$ and $\bar{X}_2$ as the random variables that take stationary distributions (suppose they exist) of processes X_1 and X_2. We have*

$$\bar{X}_2 \succeq_1 \bar{X}_1. \tag{A.1}$$

Proof of Lemma 3. Denote the *P.M.F.* (C.M.F.) of $\bar{X}_1$ and $\bar{X}_2$ as $f_1(\cdot)$ and $f_2(\cdot)$ ($F_1(\cdot)$ and $F_2(\cdot)$), respectively. We show that $F_1(x) \geq F_2(x)$ for all $x \geq 0$.

First, note that we have

$$\frac{f_1(x)}{f_2(x)} = \frac{f_1(0)}{f_2(0)} \Pi_{i=1}^{x} \frac{\mu_2(x)}{\mu_1(x)}, \tag{A.2}$$

since both X_1 and X_2 are Birth-Death processes. Since we have $\mu_1(x) \geq \mu_2(x)$ for all states $x \geq 1$, there is $\Pi_{i=1}^{x} \frac{\mu_2(x)}{\mu_1(x)} \leq 1$, which implies $f_1(0) \geq f_2(0)$ (as $f_1(\cdot)$ and $f_2(\cdot)$ are well-defined P.M.F.)

Next, we prove the desired result by contradiction. Suppose there exists some $x \geq 1$ such that $F_1(x) < F_2(x)$ and let $k := \min\{x \,|\, F_1(x) < F_2(x)\}$. By the definition of k, we have $F_1(k) < F_2(k)$ and $F_1(k-1) \geq F_2(k-1)$. These two inequalities imply that $f_1(k) < f_2(k)$.

Note by (A.2), we have that

$$\frac{f_1(x+1)}{f_2(x+1)} = \frac{f_1(x)\mu_2(x+1)}{f_2(x)\mu_1(x+1)},$$

which implies that $f_1(x) < f_2(x)$ for all $x > k$ since $\frac{\mu_2(x+1)}{\mu_1(x+1)} \leq 1$ for all $x \geq 1$. Therefore, since both $f_1(\cdot)$ and $f_2(\cdot)$ are well-defined P.M.F., we have the following inequalities

$$\sum_{i=k+1}^{\infty} f_1(i) < \sum_{i=k+1}^{\infty} f_2(i), \text{ and } \sum_{i=0}^{k} f_1(i) < \sum_{i=0}^{k} f_2(i).$$

By adding up the two inequalities above, we reach contradiction as $1 < 1$. Therefore, we have $F_1(x) \geq F_2(x)$ for all $x \geq 0$.

□

Proof of Proposition 1. The proof of the first statement follows from the result of Lemma 3 directly.

In order to prove the second statement, we show that $h_\epsilon(x,\theta)$ is a non-decreasing function of $x \geq 0$ if $0 \leq \epsilon \leq \theta \leq 1$. Then the desired result follows by the property of first order stochastic dominance.

Fix $0 \leq \epsilon \leq \theta \leq 1$ and define function,

$$\psi(x) := \frac{h_\epsilon(x,\theta)}{c} = \epsilon - \frac{1}{1+x} + \frac{[1-\epsilon(1+x)+\theta x](1-\theta)^x}{1+x}, \quad x \geq 0. \tag{A.3}$$

Furthermore, there is

$$\begin{aligned}\frac{d\psi(x)}{dx} &= \frac{1}{(1+x)^2}\{1-(1-\theta)^{1+x}+(1+x)(1-\theta)^x[1-\epsilon(1+x)+x\theta]\ln(1-\theta)\} \\ &\geq \frac{1}{(1+x)^2}\{1-[1+(1+x)\ln(1-\theta)](1-\theta)^{1+x}\},\end{aligned} \tag{A.4}$$

where the inequity follows $1-\epsilon(1+x)+x\theta \geq 1-\theta$ since $\epsilon \leq \theta$. Note we have the first order derivatives of (A.4) *w.r.t.* θ

$$-(1-\theta)^x \ln(1-\theta) \geq 0,$$

since $\theta < 1$. Therefore, we have

$$\frac{d\psi(x)}{dx} \geq \frac{1}{(1+x)^2}\{1-[1+(1+x)\ln(1-\theta)](1-\theta)^{1+x}\} \geq 0, \tag{A.5}$$

where we use $\theta = 0$ in the second inequality. Thus, we have shown that $\psi(x)$ is a non-decreasing function, so is $h_\epsilon(x,\theta)$ *w.r.t.* x. Lastly, the inequalities in (4.5) of Proposition 1 follows from the property of first order stochastic dominance and the fact function $h_\epsilon(x,\theta)$ is non-decreasing in x.

□

Proof of Proposition 2. Note that we have

$$\frac{\mathbb{E}h_\epsilon(X(\epsilon\rho),\epsilon\rho)-\mathbb{E}h_\epsilon(Y(\epsilon\rho),\epsilon\rho)}{\mathbb{E}h_\epsilon(X(\epsilon\rho),\epsilon\rho)} \leq \frac{\mathbb{E}h_\epsilon(Y(0),\epsilon\rho)-\mathbb{E}h_\epsilon(Y(\epsilon\rho),\epsilon\rho)}{\mathbb{E}h_\epsilon(Y(\epsilon\rho),\epsilon\rho)}, \tag{A.6}$$

where the inequality follows Proposition 1 as $\mathbb{E}h_\epsilon(Y(\epsilon\rho),\epsilon\rho) \le \mathbb{E}h_\epsilon(X(\epsilon\rho),\epsilon\rho) \le \mathbb{E}h_\epsilon(Y(0),\epsilon\rho)$. Furthermore, the right-hand-side of (A.6) has closed-form expression since both $Y(0)$ and $Y(\epsilon\rho)$ are Poisson random variables. That is, we have

$$\frac{\mathbb{E}h_\epsilon(Y(\epsilon\rho),\epsilon\rho) - \mathbb{E}h_\epsilon(Y(0),\epsilon\rho)}{\mathbb{E}h_\epsilon(Y(\epsilon\rho),\epsilon\rho)} = \frac{\epsilon\lambda_L - \gamma + (\gamma - \epsilon\lambda(1-\rho))\exp\left(-\frac{\lambda_L\epsilon\rho}{\gamma}\right)}{\epsilon\lambda_L - \gamma - \lambda_S\epsilon\rho + (\gamma - \lambda_L\epsilon + \epsilon\rho(\lambda_L+\lambda_S))\exp\left(-\frac{\lambda_L\epsilon\rho}{\gamma+\lambda_S\epsilon\rho}\right)} - 1$$
$$= \frac{\lambda\rho}{\gamma}\epsilon + o(\epsilon),$$

where the equality follows a first order Taylor expansion around $\epsilon = 0$. According to the definition of "*little-o*" notations, we have the desired result. □

Proof of Lemma 2. Consider a Poisson random variable $Y(\epsilon\rho)$ with load factor $\dfrac{\lambda_L}{\gamma+\lambda_S\epsilon\rho}$. We have

$$\mathbb{E}h_\epsilon(Y(\epsilon\rho),\epsilon\rho) = \frac{1}{\lambda_L}\left[\epsilon(\lambda_L - \lambda_S\rho) - \gamma + [\gamma + \epsilon(\rho(\lambda_L+\lambda_S) - \lambda_L)]\exp\left(-\frac{\epsilon\lambda_L\rho}{\epsilon\lambda_S\rho+\gamma}\right)\right]. \tag{A.7}$$

Next, we perform a third order Taylor expansions on (A.7). Define

$$J(\rho,\epsilon) := \lambda_L\left[\frac{\rho(2-\rho)}{2\gamma}\epsilon^2 - \frac{\rho^2[3\lambda_S(2-\rho)+\lambda_L(3-2\rho)]}{6\gamma^2}\epsilon^3\right], \tag{A.8}$$

so that there is

$$\mathbb{E}h_\epsilon(Y(\epsilon\rho),\epsilon\rho) = J(\rho,\epsilon) + o(\epsilon^3),$$

which implies the result according to the definition of "*little-o*" notations.

Note that one can also perform Taylor expansion over function h_ϵ around $\epsilon = 0$ directly before taking the expectation *w.r.t.* the Poisson random variable $Y(\epsilon\rho)$. However, the "higher-order" terms ($o(\epsilon^3)$) contain the random variable $Y(\epsilon\rho)$, which can be infinity. Then one need to verify that these "higher-order" terms are indeed going to zero as $\epsilon \to 0$, which can be shown using the "light-tail" property of Poisson random variables. We omit details for this procedure.

□

Proof of Theorem 1 and Corollary 1. Consider

$$\epsilon < \frac{\gamma}{\lambda_s+\lambda_L}. \tag{A.9}$$

The second order derivatives of function $J(\rho,\epsilon)$ *w.r.t.* ρ is

$$\frac{d^2J(\rho,\epsilon)}{d\rho^2} = \lambda_L\epsilon^2\left[-\frac{1}{\gamma} + \frac{\lambda_L(2\rho-1)+\lambda_S(3\rho-2)}{\gamma^2}\epsilon\right], \tag{A.10}$$

which is strictly negative according to (A.9). Thus, function $J(\rho,\epsilon)$ is strictly concave in ρ when $\epsilon < \dfrac{\gamma}{\lambda_s+\lambda_L}$ so it has a unique maximizer $\rho^*_\epsilon \in [0,1]$. By solving $\dfrac{\partial J(\rho,\epsilon)}{\partial\rho} = 0$, we have

$$\rho^*_\epsilon = \frac{\epsilon(\lambda_L+2\lambda_S) + \gamma - \sqrt{\epsilon^2(\lambda_L+2\lambda_S)^2 - 2\epsilon(\lambda_L+\lambda_S)\gamma + \gamma^2}}{\epsilon(2\lambda_L+3\lambda_S)}. \tag{A.11}$$

Next, consider

$$\hat{\rho}_\epsilon = 1 - \frac{\lambda_S}{2\gamma}$$

in (4.9). We perform a first order Taylor expansion on ρ^*_ϵ around $\epsilon = 0$ and there is

$$\rho^*_\epsilon = \hat{\rho}_\epsilon + o(\epsilon).$$

In order to prove the second statement in Theorem 1, we first show that when $\rho \in [0,1]$, function $\mathbb{E}h_\epsilon(Y(0), \epsilon\rho)$ is non-decreasing in ρ by checking the first order derivatives *w.r.t.* ρ:

$$\frac{d\,\mathbb{E}h_\epsilon(Y(0),\epsilon\rho)}{d\,\rho} = \frac{\lambda_L c(1-\rho)\exp\left(-\frac{\lambda_L \epsilon\rho}{\gamma}\right)}{\gamma}\epsilon^2 \geq 0.$$

Thus, we have

$$\mathbb{E}h_\epsilon(Y(0),\epsilon\rho) \leq \mathbb{E}h_\epsilon(Y(0),\epsilon), \quad \forall \rho \in [0,1]. \tag{A.12}$$

Next, consider a more general price $\rho(\alpha)$ that is a linear combination of $\hat{\rho}_\epsilon$ in (4.11) and 1. For $\alpha \in [0,1]$, we have

$$\begin{aligned}
&\frac{\mathbb{E}h_\epsilon(X(\epsilon\rho^*),\epsilon\rho^*) - \mathbb{E}h_\epsilon(X(\epsilon\rho(\alpha)),\epsilon\rho(\alpha))}{\mathbb{E}h_\epsilon(X(\epsilon\rho^*),\epsilon\rho^*)} \\
\leq\ &\frac{\mathbb{E}h_\epsilon(X(\epsilon\rho^*),\epsilon\rho^*) - \mathbb{E}h_\epsilon(X(\epsilon\rho(\alpha)),\epsilon\rho(\alpha))}{\mathbb{E}h_\epsilon(X(\epsilon\rho(\alpha)),\epsilon\rho(\alpha))} \\
\leq\ &\frac{\mathbb{E}h_\epsilon(Y(0),\epsilon\rho^*) - \mathbb{E}h_\epsilon(Y(\epsilon\rho(\alpha)),\epsilon\rho(\alpha))}{\mathbb{E}h_\epsilon(Y(\epsilon\rho(\alpha)),\epsilon\rho(\alpha))} \\
\leq\ &\frac{\mathbb{E}h_\epsilon(Y(0),\epsilon) - \mathbb{E}h_\epsilon(Y(\epsilon\rho(\alpha)),\epsilon\rho(\alpha))}{\mathbb{E}h_\epsilon(Y(\epsilon\rho(\alpha)),\epsilon\rho(\alpha))} \\
=\ &\frac{\lambda_L\epsilon - \gamma\left[1 - \exp\left(-\frac{\lambda_L\epsilon}{\gamma}\right)\right]}{\lambda_L\epsilon - \gamma - \lambda_S\epsilon\rho(\alpha) + ((\lambda_L+\lambda_S)\epsilon\rho(\alpha) + \gamma - \lambda_L\epsilon)\exp\left(-\frac{\lambda_L\epsilon\rho(\alpha)}{\gamma+\lambda_S\epsilon\rho(\alpha)}\right)} - 1 \\
=\ &\frac{\lambda_S}{\gamma} + o(\epsilon),
\end{aligned}$$

where the first inequality follows the definition of ρ^* and optimality; the second inequality follows Proposition 1; the third inequality follows (A.12), and the last equality follows a first order Taylor expansion around $\epsilon = 0$. This completes the proof of Corollary 1, which implies the second statement of Theorem 1.

□

Proof of Proposition 3. Before going into the proof, we define the following notations. We use $X_{\hat{\theta},\theta} = \{X_{\hat{\theta},\theta}(t), t \geq 0\}$ to denote the process describing the total number of *other* long-lived players on the platform from a focal player's perspective. Moreover, denote $N = \{N^L, N^S, N^\gamma\}$ as a three-dimensional counting process describing the total number of arrival of long-/short-lived players and departure of long-lived players, respectively. Furthermore, we split $N^S(t) = N^{Sm}(t) + N^{Sn}(t)$ representing the total numbers of short-lived arrivals that are matched and not matched by time t, respectively. Thus, we have

$$X_{\hat{\theta},\theta}(t) = X_{\hat{\theta},\theta}(0) + N^L(t) - N^\gamma(t) - N^{Sm}(t), \quad \forall t \geq 0. \tag{A.13}$$

Therefore, we can write the focal player's utility function as an integral over time:

$$V(x,\hat{\theta},\theta,\rho) = \begin{cases} \mathbb{E}^N\left[\int_0^\infty e^{-\gamma t}\mathcal{A}(X_{\hat{\theta},\theta}(t),\hat{\theta},\hat{\theta},\rho)dN^S(t)\,\middle|\,X_{\hat{\theta},\theta}(0)=x\right], & \text{if } \hat{\theta}\le\theta, \\ \mathbb{E}^N\left[\int_0^\infty e^{-\gamma t}\mathcal{A}(X_{\hat{\theta},\theta}(t),\hat{\theta},\theta,\rho)dN^S(t)\,\middle|\,X_{\hat{\theta},\theta}(0)=x\right], & \text{if } \hat{\theta}>\theta, \end{cases} \tag{A.14}$$

where function $\mathcal{A}$ defined in (5.2) represents the expected utility of the focal player upon arrival of a short-lived player. Intuitively, $e^{-\gamma t}$ comes from the P.D.F. of the exponential distribution for reneging. It is equivalent to a discount factor.

In the following proofs of the two statements in Proposition 3, we only show the result for $\hat{\theta}\le\theta$ as the counterpart follows the exact same steps.

(i) First we show that function $\mathcal{A}(x,\hat{\theta},\hat{\theta},\rho)$ is decreasing *w.r.t.* x. Note that we have

$$\frac{\partial^2\mathcal{A}(x,\hat{\theta},\hat{\theta},\rho)}{\partial x\,\partial\hat{\theta}} = -\epsilon(\hat{\theta}-\epsilon\rho)(1-\hat{\theta})^x\ln(1-\hat{\theta})\le 0,$$

according to (5.2) and $\hat{\theta}\le\epsilon\rho\le 1$. Thus, we have

$$\frac{\partial\mathcal{A}(x,\hat{\theta},\hat{\theta},\rho)}{\partial x}\le\frac{\partial\mathcal{A}(x,0,0,\rho)}{\partial x}=0.$$

Therefore, for any $x\ge 1$, there is

$$\mathcal{A}(x,\hat{\theta},\hat{\theta},\rho)\le\mathcal{A}(0,\hat{\theta},\hat{\theta},\rho). \tag{A.15}$$

According to the expression of the focal player's utility function as an integral in (A.14), we have

$$\begin{aligned} V(x,\hat{\theta},\theta,\rho) &= \mathbb{E}^N\left[\int_0^\infty e^{-\gamma t}\mathcal{A}(x,\hat{\theta},\hat{\theta},\rho)dN^S(t)\right] \\ &\le \mathbb{E}^N\left[\int_0^\infty e^{-\gamma t}\mathcal{A}(0,\hat{\theta},\hat{\theta},\rho)dN^S(t)\right] \\ &\le \mathcal{A}(0,\hat{\theta},\hat{\theta},\rho)\int_0^\infty e^{-\gamma t}\,dt = \frac{\mathcal{A}(0,\hat{\theta},\hat{\theta},\rho)}{\gamma}, \quad \forall x\ge 0, \end{aligned} \tag{A.16}$$

where the first inequality follows (A.15) and the second inequality follows the definition of N^S. Thus, there exists an upper bound $B := \dfrac{\mathcal{A}(0,\hat{\theta},\hat{\theta},\rho)}{\gamma}$ such that $V(x,\hat{\theta},\theta,\rho)\le B$ for all $x\ge 0$.

(ii) Consider a pure birth process $\tilde{X}=\{\tilde{X}(t),t\ge 0\}$ with arrival rate λ_L. Denote $\tau=\min\{t\ge 0\,|\,X_{\hat{\theta},\theta}(t)=\bar{x}\}$ and $\tilde{\tau}=\min\{t\ge 0\,|\,\tilde{X}(t)=\bar{x}\}$, representing the first times the number of players reach $\bar{x}$ under the two processes, respectively. Note that $\tilde{\tau}$ follows Erlang-k distribution with rate λ_L and $k=\bar{x}-\tilde{X}(0)$ since $\tilde{X}$ is a pure Birth (Poisson) process.

We show that if the two aforementioned processes have $X_{\hat{\theta},\theta}(0)=\tilde{X}(0)\le\bar{x}$, then there is $\mathbb{P}(\tilde{\tau}<t)\ge\mathbb{P}(\tau<t)$ for all $t\ge 0$. We show this result by coupling. As $\tilde{X}$ is a counting process, define left-continuous jump process Y such that

$$Y(t)=\tilde{X}(t)-Z(Y(t_-)),\quad t\ge 0, \tag{A.17}$$

where $Z(Y)=\{Z(Y(t_-))\,|\,t\ge 0\}$ is also a counting process with arrival rate $\gamma y+\lambda_S\mathcal{B}(y,\hat{\theta},\theta)$ when $Y(t_-)=y$ and function $\mathcal{B}$ defined in (5.4). Denote $\hat{\tau}=\min\{t\ge 0\,|\,Y(t)=\bar{x}\}$. Thus, by construction, $Y(t)\overset{D}{=}X(t)$ (equal in distribution), which implies that

$$\hat{\tau}\overset{D}{=}\tau. \tag{A.18}$$

Since $\gamma > 0$, we have $\tilde{X}(t) \geq Y(t)$ almost surely for all $t \geq 0$, which implies that

$$\tilde{\tau} \leq \hat{\tau}, \quad a.s., \tag{A.19}$$

which gives

$$\mathbb{P}(\tilde{\tau} < t) \geq \mathbb{P}(\hat{\tau} < t) = \mathbb{P}(\tau < t), \quad \forall t \geq 0, \tag{A.20}$$

where the inequality follows (A.19) and the equality follows (A.18).

By the definition of first order stochastic dominance and the fact that $e^{-\gamma t}$ is strictly decreasing *w.r.t.* t, we reach

$$\mathbb{E}_{\tau}\left[e^{-\gamma\tau}\right] \leq \mathbb{E}_{\tilde{\tau}}\left[e^{-\gamma\tilde{\tau}}\right]. \tag{A.21}$$

Now, we can write the utility functions $V(x)$, $V_{\bar{x}}(x)$ as an integrals over time similar to (A.14). For $X_{\hat{\theta},\theta}(0) = x$, there are

$$V(x,\hat{\theta},\theta,\rho) = \mathbb{E}^N\left[\int_0^{\tau} e^{-\gamma t}\mathcal{A}(X_{\hat{\theta},\theta}(t),\hat{\theta},\hat{\theta},\rho)dN^S(t) + e^{-\gamma\tau}V(\bar{x},\hat{\theta},\theta,\rho)\right], \tag{A.22}$$

$$V_{\bar{x}}(x,\hat{\theta},\theta,\rho) = \mathbb{E}^N\left[\int_0^{\tau} e^{-\gamma t}\mathcal{A}(X_{\hat{\theta},\theta}(t),\hat{\theta},\hat{\theta},\rho)dN^S(t) + 0\right]. \tag{A.23}$$

Therefore,

$$V(x,\hat{\theta},\theta,\rho) - V_{\bar{x}}(x,\hat{\theta},\theta,\rho) = \mathbb{E}_{\tau}\left[e^{-\gamma\tau}\right]V(\bar{x},\hat{\theta},\theta,\rho) \leq \mathbb{E}_{\tilde{\tau}}\left[e^{-\gamma\tilde{\tau}}\right]V(\bar{x},\hat{\theta},\theta,\rho) \leq \mathbb{E}_{\tilde{\tau}}\left[e^{-\gamma\tilde{\tau}}\right]B = B\left(1+\frac{\gamma}{\lambda_L}\right)^{-(\bar{x}-x)},$$

where $B = \dfrac{\mathcal{A}(0,\hat{\theta},\hat{\theta},\rho)}{\gamma}$, the first inequality follows from (A.21), second inequality follows from part(i) and last equality follows the moment generating function of Erlang random variables. This completes the proof.

□

B. Additional Proofs and Derivations

B.1. Existence of the stationary distribution.

It is well understood that a Birth-Death process has stationary distribution if and only if

$$\sum_{x=1}^{\infty}\Pi_{k=1}^{x}\frac{\lambda(k-1)}{\mu(k)} \leq \infty,$$

where $\lambda(x)$ is the system birth rate and $\mu(x)$ is the system death rate when system state is x. For the Birth-Death process in this paper, we have

$$\sum_{x=1}^{\infty}\frac{\lambda_L^x}{\Pi_{k=1}^{x}(\lambda_S p_k(\theta)+\gamma k)} \leq \sum_{x=1}^{\infty}\frac{\lambda_L^x}{\Pi_{k=1}^{x}\gamma k} = \sum_{x=1}^{\infty}\left(\frac{\lambda_L^x}{\gamma}\right)^k\frac{1}{x!} = \exp\left(\frac{\lambda_L}{\gamma}\right) - 1 < \infty, \tag{B.1}$$

where the last equality follows the *p.m.f.* of a Poisson random variable with load λ_L/γ. Thus, the stationary distribution of this Birth-Death process always exists for positive and finite λ_L and γ.

B.2. Myopic matching policy in Section 4.

We provide a proof of a stronger statement that implies $\rho_M = 1$ maximizes the matching rate defined in (4.3), Section 4.

LEMMA 4. *The matching rate function*

$$R(\theta) = \lambda_S \mathbb{E}_{X(\theta)}[1-(1-\theta)^{X(\theta)}],$$

defined in (3.10) *is increasing with respect to* $\theta \in [0,1]$.

Proof of Lemma 4. We first show that

$$\lambda_L = \lambda_S \mathbb{E}[p_{X(\theta)}(\theta)] + \gamma \mathbb{E}[X(\theta)]. \tag{B.2}$$

Recall the random variable $X(\theta)$ that follows the distribution function in (3.7). We have that

$$\begin{aligned}
\lambda_S \mathbb{E}[p_{X(\theta)}(\theta)] + \gamma \mathbb{E}[X(\theta)] &= \sum_{x=1}^{\infty} \frac{{\lambda_L}^x f_\theta(0)}{\Pi_{k=1}^x(\lambda_S p_k(\theta)+\gamma k)} \lambda_S p_x(\theta) + \sum_{x=1}^{\infty} \frac{{\lambda_L}^x f_\theta(0)}{\Pi_{k=1}^x(\lambda_S p_k(\theta)+\gamma k)} \gamma x \\
&= \sum_{x=1}^{\infty} (\lambda_S p_x(\theta)+\gamma x) \frac{{\lambda_L}^x f_\theta(0)}{\Pi_{k=1}^x(\lambda_S p_k(\theta)+\gamma k)} \\
&= \lambda_L \left(f_\theta(0) + \sum_{x=2}^{\infty} \frac{{\lambda_L}^{x-1} f_\theta(0)}{\Pi_{k=1}^{x-1}(\lambda_S p_k(\theta)+\gamma k)} \right) \\
&= \lambda_L \left(f_\theta(0) + \sum_{y=1}^{\infty} \frac{{\lambda_L}^{y} f_\theta(0)}{\Pi_{k=1}^{y}(\lambda_S p_k(\theta)+\gamma k)} \right) \\
&= \lambda_L \sum_{y=0}^{\infty} f_\theta(y) = \lambda_L.
\end{aligned}$$

Thus, we reach the equation in (B.2).

Next, consider $0 \le \theta_1 \le \theta_2 \le 1$. We have $X(\theta_2) \succeq_1 X(\theta_1)$ from Lemma 3 since $p_x(\theta_1) \le p_x(\theta_2)$ for all $x \in \mathbb{Z}_+$. Since function $p_x(\theta)$ is increasing in $\theta \in [0,1]$, by applying the property of first order stochastic dominance, we reach that

$$\mathbb{E}[X(\theta_2)] \le \mathbb{E}[X(\theta_1)],$$

which, according to (B.2), implies that

$$\mathbb{E}[p_{X(\theta_1)}(\theta_1)] \le \mathbb{E}[p_{X(\theta_2)}(\theta_2)].$$

This completes the proof.

□

B.3. Long-Liver players' utility function in Section 5.

We can write out the focal player's expected utility functions recursively in a heuristic manner. Fix $\hat{\theta} \le \theta$ and consider an infinitesimal time period $[t, t+\delta)$,

$$\begin{aligned}
V(x) = (1-\gamma\delta)\Big\{ &\lambda_L \delta V(x+1) + x\gamma\delta V(x-1) + (1-\lambda_L\delta - \lambda_S\delta - x\gamma\delta)V(x) \\
&+ \lambda_S \delta \left[\mathcal{A}(x,\hat{\theta},\hat{\theta},\rho) + \mathcal{B}(x,\hat{\theta},\theta)V(x-1) + \mathcal{C}(x,\hat{\theta},\theta)V(x) \right] \Big\}
\end{aligned}$$

$$= (1-\gamma\delta)\Big\{\lambda_S\delta\mathcal{A}(x,\hat{\theta},\hat{\theta},\rho) + \lambda_L\delta V(x+1) + \Big[1-\lambda_L\delta - \lambda_S\delta[1-\mathcal{C}(x,\hat{\theta},\theta)] - x\gamma\delta\Big]V(x) + \Big[\lambda_S\mathcal{B}(x,\hat{\theta},\theta) + x\gamma\Big]\delta V(x-1)\Big\}. \tag{B.3}$$

By dividing both sides with δ and then take $\delta \to 0$, we reach,

$$V(x) = \frac{\lambda_S\mathcal{A}(x,\hat{\theta},\hat{\theta},\rho) + \Big[\lambda_S\mathcal{B}(x,\hat{\theta},\theta) + x\gamma\Big]V(x-1) + \lambda_L V(x+1)}{\lambda_L + \lambda_S(1-\mathcal{C}(x,\hat{\theta},\theta)) + (x+1)\gamma}.$$

The derivation for $\hat{\theta} > \theta$ follows the same steps and, thus, it is omitted.

B.4. Derivations for numerical procedure with heterogeneous players in Section 6.

First, we need to derive a focal player's expected utility function similar to (5.8). Note that, upon being matched with a mismatch angle ϕ, a long-lived player's utility is $P - c\phi = c(\bar{\epsilon}\bar{\rho} - \phi)$, which is not affected by the heterogeneity among short-lived players. Furthermore, if the price $\bar{\rho}$ is fixed, the only difference for long-lived players when facing short-lived players with heterogeneity is that they need to consider the market entry of short-lived players. Following the heuristic derivation in (B.3) and assume the focal player uses threshold $\hat{\theta}$ and all other long-lived players use threshold θ. Fix $\hat{\theta} \le \theta$, we have

$$V(x) = (1-\gamma\delta)\Big\{\lambda_L\delta V(x+1) + x\gamma\delta V(x-1) + (1-\lambda_L\delta - \lambda_S\frac{\sigma+1-\bar{\rho}}{2\sigma}\delta - x\gamma\delta)V(x) + \lambda_S\frac{\sigma+1-\bar{\rho}}{2\sigma}\delta\Big[\mathcal{A}(x,\hat{\theta},\hat{\theta},\bar{\rho}) + \mathcal{B}(x,\hat{\theta},\theta)V(x-1) + \mathcal{C}(x,\hat{\theta},\theta)V(x)\Big]\Big\}$$

which by dividing both sides with δ and then take $\delta \to 0$, leads to

$$V(x) = \frac{\bar{\lambda}_S\mathcal{A}(x,\hat{\theta},\hat{\theta},\bar{\rho}) + \Big[\bar{\lambda}_S\mathcal{B}(x,\hat{\theta},\theta) + x\gamma\Big]V(x-1) + \lambda_L V(x+1)}{\lambda_L + \bar{\lambda}_S(1-\mathcal{C}(x,\hat{\theta},\theta)) + (x+1)\gamma},$$

where

$$\bar{\lambda}_S = \lambda_S\frac{\sigma+1-\bar{\rho}}{2\sigma}. \tag{B.4}$$

The derivation of value function when $\hat{\theta} > \theta$ follows the same procedure so we omit it. In summary, a long-lived player's value function V solves

$$V(x) = \begin{cases} \dfrac{\bar{\lambda}_S\mathcal{A}(x,\hat{\theta},\hat{\theta},\bar{\rho}) + \Big[\bar{\lambda}_S\mathcal{B}(x,\hat{\theta},\theta) + x\gamma\Big]V(x-1) + \lambda_L V(x+1)}{\lambda_L + \bar{\lambda}_S(1-\mathcal{C}(x,\hat{\theta},\theta)) + (x+1)\gamma}, & \text{if } \hat{\theta} \le \theta, \\ \dfrac{\bar{\lambda}_S\mathcal{A}(x,\hat{\theta},\theta,\bar{\rho}) + \big[\bar{\lambda}_S\mathcal{B}(x,\theta,\theta) + x\gamma\big]V(x-1) + \lambda_L V(x+1)}{\lambda_L + \bar{\lambda}_S(1-\mathcal{C}(x,\hat{\theta},\theta)) + (x+1)\gamma}, & \text{if } \hat{\theta} > \theta, \end{cases} \tag{B.5}$$

with boundary condition

$$\bar{\lambda}_S\mathcal{A}(0,\hat{\theta},\theta,\bar{\rho}) = (\lambda_L + \bar{\lambda}_S\mathcal{C}(0,\hat{\theta},\theta) + \gamma)V(0) - \lambda_L V(1). \tag{B.6}$$

As Proposition 3 is independent of the choice of short-lived players' arrival rate, it still applies here. Therefore, fix $\bar{x} > 1$ and define function $V_{\bar{x}}(x,\hat{\theta},\theta,\bar{\rho})$ as the solution to the system of equations that solves (B.6) and (B.5) for $x \in \{0,...,\bar{x}-1\}$ with $V_{\bar{x}}(\bar{x},\hat{\theta},\theta,\bar{\rho}) = 0$. In our numerical procedure, for any $\bar{\rho} \in [0, 1+\sigma]$, we compute the equilibrium matching threshold of long-lived player in the same way as in (5.13):

$$\bar{\theta}_L \in \operatorname*{arg\,max}_{\hat{\theta}\in[0,\bar{\epsilon}\bar{\rho}]} \mathbb{E}[V_{\bar{x}}(X(\bar{\theta}_L),\hat{\theta},\bar{\theta}_L,\bar{\rho})].$$