Revenue from matching platforms

Philip Marx
Economics Department, Louisiana State University

James Schummer
MEDS Department, Kellogg School of Management, Northwestern University

We consider the pricing problem of a platform that matches heterogeneous agents using match-contingent fees. Absent prices, agents on the short side of such markets capture relatively greater surplus than those on the long side (Ashlagi et al. 2017). Nevertheless we show that the platform need not bias its price allocation toward either side. With independently drawn preferences, optimal price allocation decisions are independent of market size or imbalance; furthermore, changes in the optimal price level move both sides' prices in the same direction. In contrast, preference homogeneity biases price allocation in a direction that depends on the form of homogeneity; furthermore, changes in market imbalance move both sides' prices in opposite directions. These effects arise due to the exclusivity of matchings in two-sided market settings.

Keywords. Pricing, matching, platforms.

JEL classification. C78, D40.

1. Introduction

The proliferation of online, "winner-take-all" matching platforms has led to intensified interest in the study of platform pricing. While an established literature on two-sided markets explains much about such pricing, it has mostly[1] abstracted away from two real-world characteristics of markets that can lead to market size effects on pricing that otherwise would not be present. Specifically, we consider settings where *exclusive* (one-to-one) partnerships are created among *horizontally differentiated* agents.

(I) Exclusivity. Canonical models of two-sided markets realistically represent certain platforms (video game consoles, credit cards, newspapers, etc.) by assuming that each participating agent interacts with *all* (or a constant fraction of) agents on the other side of the market. Our interest is in platforms that exist specifically to create *one-to-one* (or capacity constrained) matchings (e.g., ride-sharing or matchmaking services).

[1]Exceptions are discussed in Section 1.2.

(II) Heterogeneity.   Many models of two-sided interactions consider some form of homogeneity in preferences; e.g., that agents commonly rank potential partners or express indifference over partners. Such assumptions are realistic in certain environments, but in others—especially those exhibiting exclusivity—agents' preferences over potential partners are heterogeneous; examples include dating or headhunting services.

Characteristics (I) and (II) lead to a third distinction between existing work and ours: consideration of the matching process itself. Agents care not only about the price they pay to the platform, but about the identity of the partner with whom they are matched.[2] Though the platform controls prices, it may or may not control the matching process itself. Centralized platforms (e.g., ride sharing) may be able to directly specify matching outcomes. Decentralized ones (e.g., dating sites) do so only indirectly by setting the rules by which agents interact to make pairing decisions. This leaves the question of *which* matching outcomes might actually occur on such platforms.

When search and information frictions are low, it is reasonable to assume that decentralized platforms yield matching outcomes in the core, i.e., stable matchings à la Gale and Lloyd (1962). Indeed, Adachi (2003) obtains such outcomes in the equilibria of a decentralized matching model with random encounters. Hitch et al. (2010) validate this notion empirically, demonstrating that interactions at an online dating site resemble stable matching outcomes with respect to estimated preferences.

Though stability is a natural assumption, our main results apply to any platform whose matching outcomes satisfy a weaker set of assumptions. Specifically, we consider any platform (decentralized or not) whose matching outcomes are (a) individually rational and (b) sensitive only to the agents' ordinal preferences over compatible partners. Examples include platforms yielding stable matchings, serially dictatorial matchings, or maximal matchings subject to individual rationality. Our "black box" treatment of the matching process allows us to bypass assumptions on the agents' behavior and information, while also clarifying the intuition underlying our results.

1.1 *Overview*

Our work sits between two established literatures. A *two-sided markets* literature, as described above, analyzes pricing in environments that differ from ours. A *two-sided matching* literature pioneered by Gale and Lloyd (1962) considers environments satisfying (I) and (II), but sets aside pricing. We ask how pricing structure in the latter kinds of environments may differ from those studied in the former, particularly due to the presence of market size imbalance and of preference homogeneity (or correlation).

The resulting insights depart from those obtained from either literature. For instance, when matchings are *non*exclusive it is primarily the *marginal* distribution of agents' values that determines the platform's transaction fees, independently of market imbalance and preference correlation. Intuitively, if an agent is not crowded out of a potential pairing by that partner's interaction with others, then revenue-maximizing transaction fees for that pair can be calculated regardless of the presence of other agents.

---

[2]In contrast, without capacity constraints on interactions, it is natural simply to assume that a pair of agents interacts whenever it is mutually beneficial, given prices.

When matchings are *exclusive*, however, the platform's pricing decision can be affected by both absolute and relative market size in a way that depends on the structure of preference correlation (homogeneity). Generally speaking, market size drives the sum of prices, or *price level* (Section 4), while market imbalance and preference homogeneity interact to drive relative prices across the two sides, or *price allocation* (Section 5).

To establish that the latter effect requires this interaction, we first consider heterogeneous (independently drawn) preferences (Section 3). Here we show that neither market size nor market imbalance affects the platform's price allocation decision. This conclusion is not only a baseline from which to consider homogeneous preferences, but also establishes another departure from intuition established in previous work. Ashlagi et al. (2017) show that market imbalance leads to higher normalized payoffs for "scarce" agents on the short side of stable matching markets than for agents on the long side. One might thus expect a monopolistic platform to capture these imbalanced payoffs by charging a relatively higher price to the scarce agents. We show that this is not the case when agents are charged in the form of transaction fees; market imbalance by itself does not justify imbalanced prices.

To illustrate with a simple example,[3] consider a setting with only one agent ("man") on one side of a market who can pair with at most one of ten agents ("women") on the other, each agent drawing an independent and identically distributed (i.i.d.) value from $[0, 1]$ for each potential partner. Assume that the platform yields a (single) match whenever at least one man–woman pair is *compatible*: their values for each other exceed their respective transaction fees. By setting the man's fee to be 0.5 and the women's fee to be 0, a match is created whenever at least one of the man's ten draws exceeds 0.5. Swapping the two sides' fees, a match is created whenever at least one of the ten women draws her value in excess of 0.5. Both scenarios thus yield the same expected revenue: there is no strict benefit from charging a higher price to the man.[4]

This conclusion relies on two assumptions. First is that the matching outcome depends only on each pair's compatibility. If the outcome is sensitive to additional preference information, the conclusion need not hold. We establish weak assumptions on the matching process (Section 2.1) that preserve the example's conclusion more generally. Second is preference heterogeneity: i.i.d. values imply that (i) the man views the women as "differentiated" and (ii) the women have "private values" for the man. Introducing preference homogeneity along either of these dimensions also alters this conclusion as we discuss in Section 5.

While our main focus is on price *allocation*, we consider price *level* in Section 4. Unlike the former, price level decisions depend on the specification of the matching process. Focusing on *stable* matching outcomes (motivated above), we provide an expression that yields a lower bound for the stable platform's revenue and we illustrate that thicker markets lead to higher price levels. The expression exactly describes expected revenue for a serially dictatorial matching platform.[5]

---

[3]We are grateful to a referee for suggesting this illustration.

[4]With additional assumptions, the optimal fees are equal (Proposition 1), thus preserving the surplus-imbalance result of Ashlagi et al. (2017) even net of optimal prices.

[5]A ranking between stability and a form of serial dictatorship is first established in a related model by Arnosti (2016).

1.2 *Other approaches*

Platform pricing has been studied in ways that differ from both our setting (exclusive matchings among agents with heterogeneous preferences) and our pricing method (transaction fees). Most prominently, the two-sided markets literature alluded to above (pioneered by Rochet and Tirole 2003, 2006, and Armstrong 2006) focuses on nonexclusive matchings, with homogeneity in agents' fixed or variable benefits from joining the platform; see also Weyl (2010). Besides transaction fees, access fees also have been considered, particularly where agents also value platform membership itself.

In our setting, where value is generated solely from exclusive matching, we set aside access fees to preserve ex post individual rationality (IR). If agents decide whether to join the platform before learning their preferences over the other participants, optimal access fees simply charge each side the expected surplus of joining the platform, violating IR for some agents ex post. Notably, such access fees would typically be biased toward the shorter side of the market as intuition might suggest; under stability this follows from Ashlagi et al. (2017). In contrast, we show that transaction fees induce no such bias if preferences are heterogeneous.

Related work on exclusive matching platforms analyzes other forms of up-front payments, but in a setting of homogeneous (vertical) preferences, where efficient (or stable) matchings are assortative. First, Damiano and Li (2007)[6] offer agents a *menu* of access fees, each giving entry to a "club" that randomly matches members. Knowing their types when they choose clubs, agents end up coarsely sorted by type. The authors' main concern is the inefficiency (coarseness) of revenue-maximizing menus.

Second, one can reinterpret the money-burning signalling model of Hoppe et al. (2009) as a platform that runs an all-pay auction, assortatively matching agents with respect to their bids. Their results imply that an increase in market imbalance would increase total revenue earned from the long side and increase total surplus on the short side.

A search market literature also addresses pricing and exclusive matching. Bloch and Ryder (2000) consider agents (again with vertical preferences) who have the outside option to find a partner via a search market with friction. They show that a platform using flat fees loses low-type agents to the search market, while fees proportional to surplus cause high types to defect. Analogous results in the case of transferable utility are shown by Rubinstein and Wolinsky (1987) and Yavaş (1992).

The related, wider problem of preventing defection from the platform is addressed by Spiegler (2000) through more general (contractual) pricing arrangements. Roughly speaking, his solution induces agents to stay on the platform by offering certain compensations to one side of the market whenever agents on the other side defect. He further shows that such schemes can be constructed in a way that limits the platform's liability.

Other topics beyond our scope, such as platform competition and information structure, use models that further depart from ours. Depending on the setting, platform

---

[6]Also see Damiano and Hao (2008) for an analysis of competing platforms.

competition can be softened when platforms offer nonexclusive membership (Caillaud and Jullien 2003), each platform restricts agent participation on a unique side of the market (Halaburda et al. 2018), and two platforms price to be cheaper (and larger) on opposite sides of the market (Ambrus and Argenziano 2009). In terms of information structure, Biglaiser (1993) studies platforms as "experts" that reduce adverse selection; Fershtman and Pavan (2020) and Gomes and Pavan (2016) study the impact of information structure on platform pricing and design. We abstract away from informational considerations by "black-boxing" the matching process.

2. Model

There are two finite sets of agents, referred to as men $M = \{1, 2, \dots, M\}$ and women $W = \{1, 2, \dots, W\}$. A (one-to-one) *matching* is a function $\nu\colon M \times W \to M \times W$ that satisfies the following usual conditions for all $(m, w) \in M \times W$: (i) $\nu(m) \in W \cup \{m\}$, (ii) $\nu(w) \in M \cup \{w\}$, and (iii) $\nu(m) = w$ if and only if $\nu(w) = m$. We say agent $i \in M \cup W$ is *unmatched* (or *single*) at ν when $\nu(i) = i$. We say $(m, w) \in M \times W$ is a *marriage* when $\nu(m) = w$. The *number of marriages* in ν is denoted $|\nu|$.

   If man $m \in M$ is matched to woman $w \in W$, m obtains *value* $u_m(w) \in [0, 1]$ and w obtains $u_w(m) \in [0, 1]$. The value of being unmatched is zero (denoted $u_i(i) \equiv 0$). These normalizations are not critical to our results.

   A profile of values, denoted $u = ((u_m)_M, (u_w)_W)$, is randomly drawn from a joint distribution. Heterogeneous preferences (Section 3 and Section 4) are represented by independently drawn values. Namely, a *random economy* is one where each value $u_m(w)$ is independently drawn from (marginal) distribution $F_{\mathcal{M}}$ and each $u_w(m)$ is independently drawn from $F_{\mathcal{W}}$. The corresponding densities $(f_{\mathcal{M}}, f_{\mathcal{W}})$ are continuously differentiable with positive support on $[0, 1]$. Homogeneous preferences are represented by correlated values as constructed in Section 5.

   Agents make transfers to the platform as described below. At a matching ν, an agent i who pays x to the platform has a payoff of $u_i(\nu(i)) - x$.

2.1 *Matching platforms*

A matching platform specifies prices and, given these prices, yields a matching outcome as a function of the agents' realized values. In terms of prices, we focus on platforms that charge agents via match-contingent fees. That is, an agent's payment to the platform is a function only of (i) the side of the market to which the agent belongs and (ii) whether or not the agent ends up in a marriage. Formally, *prices* are a pair $p = (p_{\mathcal{M}}, p_{\mathcal{W}}) \in [0, 1]^2$, where matched men and women pay the platform $p_{\mathcal{M}}$ and $p_{\mathcal{W}}$, respectively.[7] The payments of unmatched agents are normalized to zero.

   Fixing such prices p, a platform is defined by the manner in which it creates matchings as a function of the agents' realized values. A *matching process* μ is a function that, for any p and u, yields a matching $\nu = \mu_{p,u}\colon M \times W \to M \times W$.

---

[7]The restriction to $[0, 1]$ is innocuous since the support of values is on $[0, 1]$; e.g., see the arguments in the proof of Lemma 2.

There are multiple interpretations of the function μ, which is a "black box" representation of how agents end up being matched. At one extreme, imagine a centralized platform that can fully observe the agents' preferences, u. Such a platform may be able to fully dictate the matching outcome by choosing an arbitrary μ, e.g., one that maximizes the number of marriages. In contrast, a decentralized (or uninformed) platform may be able to specify only some set of rules by which decentralized agents interact and form marriages on the platform. In this case, $\mu_{p,u}$ could represent an equilibrium outcome of the game induced by such rules under prices p at a realization u.[8] Earlier we motivated the assumption that $\mu_{p,u}$ represents a stable matching with respect to preferences induced by p and u. Nevertheless, our main results apply to matching processes that satisfy weaker conditions.

First, we restrict attention to ex post individually rational outcomes. Given prices $p = (p_{\mathcal{M}}, p_{\mathcal{W}})$ and realized values u, we call the pair $(m, w) \in M \times W$ *p-compatible (at u)* if both $u_m(w) \geq p_{\mathcal{M}}$ and $u_w(m) \geq p_{\mathcal{W}}$. The matching process μ is *individually rational* if, for any p and u,

$$\forall (m, w) \in M \times W, \quad [\mu_{p,u}(m) = w] \implies [m \text{ and } w \text{ are } p\text{-compatible}].$$

The main assumption we make on the matching process is that its outcomes are sensitive only to the agents' ordinal preferences over compatible partners. Condition 1 requires the matching outcome to remain constant whenever a change in prices (p, p') and/or realized values (u, u') has no effect on (i) the set of compatible pairs of agents and (ii) the agents' ordinal preferences over such partners.

Condition 1. A matching process μ is *ordinal subject to compatibility* (hereafter, ordinal) if, for any p, p' and u, u', we have $\mu_{p,u} = \mu_{p',u'}$ whenever the following two conditions hold.

(i) The set of p-compatible pairs at u equals the set of p'-compatible pairs at u'.

(ii) For all such p-compatible pairs (m, w), (m, w'), and (m', w) at u,

$$u_m(w) > u_m(w') \quad \Leftrightarrow \quad u'_m(w) > u'_m(w')$$
$$u_w(m) > u_w(m') \quad \Leftrightarrow \quad u'_w(m) > u'_w(m').$$

Many natural matching processes satisfy this condition, such as those that yield man-optimal (or woman-optimal) stable matchings with respect to the agents' ordinal preferences induced by u and p. This follows from the fact that the deferred acceptance algorithm that calculates such matchings is a function only of agents' ordinal preferences over compatible partners; see Section 4. Other examples include those based on certain serially dictatorial choice functions (again see Section 4) and those that choose maximal matchings subject to individual rationality. In general, whenever the matching process represents outcomes that result from some strategic behavior of agents on

---

[8]That is, μ can represent an implementable social choice function under some solution concept. In the case of stability, see Roth (1984a), Kara and Sönmez (1996), and Alcalde (1996).

a decentralized platform, the condition conveys the idea that the agents' behavior depends only on their realized ordinal preferences over "relevant" (i.e., compatible) potential partners.

Conversely, Condition 1 rules out matching processes that are sensitive to preference intensity or to preferences over incompatible partners. Consider decentralized platforms on which pairs are created only when both partner's net payoffs exceed some positive threshold $c > 0$. In this case, a small price increase might change the outcome of the matching process without affecting ordinal preferences or compatibility, violating the condition.

Some of our price allocation results apply in the natural case that higher prices decrease the number of created marriages. In fact, we only need this *law of demand* assumption to hold in expectation.

CONDITION 2. A matching process μ is *monotonic* if the expected number of realized marriages is strictly decreasing in prices: for any pairs of prices $p \lneqq p'$, we have $E_u(|\mu_{p,u}|) > E_u(|\mu_{p',u}|)$.

All of the matching processes mentioned immediately after Condition 1 also satisfy Condition 2. In fact, they satisfy a stronger, pointwise version of monotonicity: for all u, $|\mu_{p,u}| \geq |\mu_{p',u}|$, with strict inequality for some positive density of us.[9]

Condition 1 and Condition 2 suffice to prove our main results on price allocation in Section 3. However our discussion of homogeneous preferences in Section 5 is simplified by restricting attention to "maximal" matching processes, in the sense that at least one member of any p-compatible pair must belong to some marriage.

CONDITION 3. A matching process μ is *weakly unimprovable* if, for any prices p and realized values u, there exists no p-compatible pair $(m, w) \in M \times W$ such that $\mu_{p,u}(m) = m$ and $\mu_{p,u}(w) = w$.

The processes mentioned after Condition 1 are all weakly unimprovable.

3. PRICE ALLOCATION

Fixing a matching process μ and prices $p = (p_\mathcal{M}, p_\mathcal{W})$, the platform's revenue is the price *level*, $p_T \equiv p_\mathcal{M} + p_\mathcal{W}$, multiplied by the number of created marriages, $|\mu_{p,u}|$. For a random economy, the number of created marriages is a random variable that we denote below by K_p. Expected revenue is thus $p_T \cdot E(K_p) \equiv (p_\mathcal{M} + p_\mathcal{W}) \cdot E(K_p)$.

Even if we fix a price level p_T, the *allocation* of p_T between $p_\mathcal{M}$ and $p_\mathcal{W}$ typically affects the distribution of K_p, and thus affects the platform's expected revenue.[10] Therefore, the platform's pricing decision can be viewed as a choice of a total price level p_T followed by a decision of how to allocate p_T between the two sides:

$$\max_{p=(p_\mathcal{M}, p_\mathcal{W})} (p_\mathcal{M} + p_\mathcal{W}) E(K_p) = \max_{p_T} \Big(p_T \cdot \underbrace{\max_{p_\mathcal{M} + p_\mathcal{W} = p_T} E(K_p)}_{\text{price allocation decision}} \Big).$$

---

[9]In the case of stability, this follows from a result of Gale and Sotomayor (1985).

[10]Thus our model fits Rochet and Tirole's (2006) definition of two-sided markets.

Focusing on the price allocation decision, we first show that price allocation affects the distribution of marriages K_p only to the extent that it affects the probability that any arbitrary man–woman pair is p-compatible. Define the *incompatibility parameter* $q(p)$ to be the probability that an arbitrary pair $(m, w) \in M \times W$ is incompatible at prices $p = (p_{\mathcal{M}}, p_{\mathcal{W}})$:

$$q(p) = q(p_{\mathcal{M}}, p_{\mathcal{W}}) \equiv F_{\mathcal{M}}(p_{\mathcal{M}}) + F_{\mathcal{W}}(p_{\mathcal{W}}) - F_{\mathcal{M}}(p_{\mathcal{M}})F_{\mathcal{W}}(p_{\mathcal{W}}).$$

Lemma 1 states that for ordinal matching processes, the distribution of marriages is a function only of $q(p)$: All price pairs with the same incompatibility parameter yield the same expected number of marriages.[11]

LEMMA 1 (The distribution of marriages depends only on $q(p)$). *Suppose the matching process μ is ordinal. For any $p' = (p'_{\mathcal{M}}, p'_{\mathcal{W}})$ and $p'' = (p''_{\mathcal{M}}, p''_{\mathcal{W}})$ that satisfy $q(p') = q(p'')$, $K_{p'}$ and $K_{p''}$ have the same distribution.*

PROOF. For any prices p, since values $u_i(j)$ are drawn independently, the probability that any pair (m, w) will be p-compatible is $1 - q(p)$, independently of any other pair's compatibility. Therefore, the probability that the set of p-compatible pairs will be some arbitrary set $C \subseteq M \times W$ is

$$\left(1 - q(p)\right)^{|C|} \cdot q(p)^{M \cdot W - |C|},$$

which is a function only of $q(p)$. Since $q(p') = q(p'')$, p' and p'' induce the same distribution over all possible sets of compatible pairs, $C \subseteq M \times W$.

Fix some realized set of compatible pairs, C. Since values are drawn independently, an agent's ordinal preference ranking of his/her compatible partners in C is uniformly random regardless of prices. That is, conditional on any such C, p' and p'' induce the same (uniformly random) distribution over ordinal preferences.

The outcome of an ordinal matching process depends only on the realization of (i) the set of compatible pairs C and (ii) the agents' ordinal preferences over their partners in C. Since p' and p'' induce the same distribution over (i) and (ii), they induce the same distribution over matching outcomes; $K_{p'}$ and $K_{p''}$ thus have the same distribution. □

A general implication of Lemma 1 is that the platform does not bias its price allocation toward any particular side of the market as a function of market size imbalance (i.e., the size of M relative to W). We formalize this in two ways, through Theorem 1 (when $F_{\mathcal{M}} = F_{\mathcal{W}}$) and Theorem 2.

Consider the special case in which there are no ex ante differences between the two sides of the market other than size, i.e., where $F_{\mathcal{M}} = F_{\mathcal{W}}$. The expected number of marriages (and, hence, expected revenue) is then a symmetric function of $p_{\mathcal{M}}$ and $p_{\mathcal{W}}$ regardless of any imbalance between M and W. This follows immediately from Lemma 1 and the symmetry of $q(\ )$ in its arguments, so we omit a formal proof.

---

[11]This conclusion may not hold when μ is not ordinal. See Example 1 in the Appendix.

THEOREM 1 (Revenue symmetry). *Suppose the matching process is ordinal and that $F_\mathcal{M} = F_\mathcal{W}$. Let prices $p = (p_\mathcal{M}, p_\mathcal{W})$ and $p' = (p'_\mathcal{M}, p'_\mathcal{W})$ be such that $p_\mathcal{M} = p'_\mathcal{W}$ and $p_\mathcal{W} = p'_\mathcal{M}$. Then p and p' yield the same expected revenue to the platform.*

Thus any expected revenue earned by charging a relatively higher price to the short side of the market could be achieved by reversing the price list. This contrasts with the intuition (mentioned in Section 1.1) to charge more to the side of the market with higher per capita surplus. Optimal transaction fees are not biased toward the short side of the market, even if that side obtains higher gross payoffs from participation.

In light of this, it is natural to suspect that, at least for symmetric distributions $F_\mathcal{M} = F_\mathcal{W}$, a revenue-maximizing price allocation charges both sides equal prices regardless of any market imbalance. This is indeed the case when the value distributions satisfy a standard hazard rate condition. We say that F_i has a *strictly increasing hazard rate* if $h_i(x) \equiv f_i(x)/(1 - F_i(x))$ is strictly increasing in $x \in [0, 1]$. (Omitted proofs are provided in the Appendix.)

PROPOSITION 1 (Symmetry with monotone hazard rate). *Suppose the matching process is ordinal and monotonic, and that $F_\mathcal{M} = F_\mathcal{W}$ has a strictly increasing hazard rate. For any $0 < p_T < 2$, prices $p^*_\mathcal{M} = p^*_\mathcal{W} = p_T/2$ uniquely maximize expected revenue subject to the constraint $p_\mathcal{M} + p_\mathcal{W} = p_T$.*

In particular, *unconstrained* revenue-maximizing prices satisfy $p^*_\mathcal{M} = p^*_\mathcal{W}$ whenever $F_\mathcal{M} = F_\mathcal{W}$ has monotone hazard rate. Without the hazard rate condition, the platform might strictly benefit by charging two unequal prices; see Example 2 in the Appendix. Even in this case, the platform is indifferent about which side receives the higher price (Theorem 1).

Obviously when $F_\mathcal{M} \neq F_\mathcal{W}$, the revenue-maximizing price allocation need not be symmetric. Nevertheless, it remains independent of market size since optimal price allocations are simply those that minimize the incompatibility parameter, $q(\ )$.

THEOREM 2 (Optimal price allocation is independent of market size). *Suppose the matching process is ordinal and monotonic. For any $0 < p_T < 2$ and prices $p^*_\mathcal{M} + p^*_\mathcal{W} = p_T$, the following two statements are equivalent.*

  (i) *Prices $(p^*_\mathcal{M}, p^*_\mathcal{W})$ minimize $q(p_\mathcal{M}, p_\mathcal{W})$ subject to the constraint $p_\mathcal{M} + p_\mathcal{W} = p_T$.*

  (ii) *Prices $(p^*_\mathcal{M}, p^*_\mathcal{W})$ maximize expected revenue subject to the constraint $p_\mathcal{M} + p_\mathcal{W} = p_T$.*

PROOF. Under the constraint $p_\mathcal{M} + p_\mathcal{W} = p_T > 0$, $(p^*_\mathcal{M}, p^*_\mathcal{W})$ maximizes expected revenue if and only if $(p^*_\mathcal{M}, p^*_\mathcal{W})$ maximizes the expected number of marriages. By Lemma 1, the expected number of marriages is a function only of $q(\ )$.

To see that it is a *decreasing* function of $q(\ )$, consider any $0 \leq q' < q'' \leq 1$. Clearly there exist prices $p' \lneqq p''$ with $q(p') = q' < q(p'') = q''$. By monotonicity, p'' yields lower expected marriages than p'. Maximizing expected revenue is thus equivalent to minimizing $q(\ )$. □

While the optimal *allocation* of a given price level, p_T, is independent of market sizes (M and W), the choice of p_T itself is typically affected by market sizes (see Section 4). The resulting indirect effect on price allocation is ambiguous without further assumptions. Nevertheless, under the hazard rate condition, we can show that optimally allocated prices ($p_{\mathcal{M}}^*(p_T)$, $p_{\mathcal{W}}^*(p_T)$) are nondecreasing in the price level p_T.

PROPOSITION 2 (Price co-movement). *Suppose the matching process is ordinal and monotonic, and that $F_{\mathcal{M}}$ and $F_{\mathcal{W}}$ have strictly increasing hazard rates. Then the following statements hold.*

- *For any $0 < p_T < 2$, there are unique prices, ($p_{\mathcal{M}}^*(p_T)$, $p_{\mathcal{W}}^*(p_T)$), that maximize expected revenue subject to the constraint $p_{\mathcal{M}} + p_{\mathcal{W}} = p_T$.*

- *These optimally allocated prices co-vary in price level: $p_{\mathcal{M}}^*(p_T)$ and $p_{\mathcal{W}}^*(p_T)$ are nondecreasing in p_T.*

As a corollary of this result, any arbitrary change in market sizes affects both sides' (unconstrained) revenue-maximizing prices in the same direction. That is, fix all of the primitives of our model other than market sizes. Suppose that prices ($\tilde{p}_{\mathcal{M}}^*$, $\tilde{p}_{\mathcal{W}}^*$) maximize expected revenue for market sizes ($\tilde{M}$, $\tilde{W}$), while prices ($\hat{p}_{\mathcal{M}}^*$, $\hat{p}_{\mathcal{W}}^*$) maximize expected revenue for market sizes ($\hat{M}$, $\hat{W}$). Then by Proposition 2, either ($\tilde{p}_{\mathcal{M}}^*$, $\tilde{p}_{\mathcal{W}}^*$) $\leqq$ ($\hat{p}_{\mathcal{M}}^*$, $\hat{p}_{\mathcal{W}}^*$) or ($\tilde{p}_{\mathcal{M}}^*$, $\tilde{p}_{\mathcal{W}}^*$) $\geqq$ ($\hat{p}_{\mathcal{M}}^*$, $\hat{p}_{\mathcal{W}}^*$).

We explore two important points about this conclusion. First, one cannot determine which of these two inequalities holds without specifying how the matching process varies across different market sizes. Section 4 considers price level in the natural case where μ selects stable matchings at all market sizes. Second, the co-movement of optimal prices hinges on the assumption of independently drawn preferences. Section 5 shows how different forms of preference homogeneity overturn this conclusion in different directions.

4. PRICE LEVEL

In this section, we consider the topic of optimal price level (p_T) for matching processes that yield *pairwise stable* matchings à la Gale and Lloyd (1962) as motivated in the Introduction. In our context, a matching is stable with respect to prices p when (i) each matched pair is p-compatible, and (ii) no "blocking pair" of agents can match with each other and obtain strictly higher values (net of price) than in the matching. The well known deferred acceptance (DA) algorithm of Gale and Lloyd (1962) computes such a matching, which furthermore contains the same number of marriages as any other stable matching (Roth 1984b). Therefore, we restrict attention to the following DA matching process without loss of generality.

DEFINITION 1. The *DA matching process*, μ^{DA}, yields, for any prices p and values u, the matching $\mu_{p,u}^{\text{DA}}$ that results from the following algorithm. Initialize all men to be eligible and, in each round $t = 1, 2, \ldots$, execute the following steps.

Step t.1. Each eligible man proposes to his favorite (highest-valued)[12] p-compatible woman among those to whom he has not already proposed. (If no such woman exists, he proposes to no one.)

Step t.2. Each woman is tentatively serially dictatorially matched to her favorite man among those who have proposed to her (if any exist). These men become ineligible; all others become eligible. If each eligible man has proposed to each of his p-compatible partners, the algorithm outputs the tentative matching; otherwise execute round $t + 1$.

Observe that μ^{DA} satisfies the conditions of Section 2.1. Condition 1 is satisfied since the algorithm is a function only of the agents' ordinal preferences induced by p and u. Condition 2 follows from Gale and Sotomayor (1985): a price increase induces "preference truncations," shrinking the set of matched agents. Condition 3 is immediate from stability.

To analyze price level in relation to market size and imbalance, we would like a tractable expression for expected revenue under μ^{DA} as a function of both prices $(p_{\mathcal{M}}, p_{\mathcal{W}})$ and market sizes (M, W). Due to the combinatorial nature of this problem, such an expression remains out of reach. Instead we draw conclusions by first demonstrating that this expected revenue is approximated (and bounded) by the expected revenue obtained under the following (compatibility-constrained) serially dictatorial matching process (which also satisfies the conditions in Section 2.1).

DEFINITION 2. The *serial dictatorship (SD) matching process*, μ^{SD}, yields, for any prices p and values u, the matching $\mu^{\mathrm{SD}}_{p,u}$ that results from the following algorithm. Initialize all men to be eligible and, in each round $t \in \{1, 2, \ldots, W\}$, execute the following step.

Step t. Woman $w = t$ is matched to m, her favorite p-compatible man among those who are eligible. (If no such man exists, she is single.) Man m is removed from the set of single men.

Arnosti (2016) formally shows that the expected number of stable marriages is bounded by that obtained under a similar procedure. He considers large matching markets with short, constant-length preference lists on one market side, while our preference lists are "shortened" via prices and thus not of constant length as markets grow large. Since this technical difference is minor, we merely verify Arnosti's bound in our setting through simulation (see Figure 2) rather than extend his formal arguments to our setting.

Additionaly, we provide a closed-form expression for expected revenue (or marriages) under μ^{SD}, requiring the following notation.

---

[12]We ignore the zero probability event of ties. One form of preference homogeneity in Section 5 yields ties that can be broken arbitrarily without affecting the results.

DEFINITION 3. *For any real number $q \in [0, 1]$, the q-analog of integer $j \in \mathbb{Z}$ and the q-factorial of j are, respectively, defined as*

$$[j]_q \equiv 1 + q + \cdots + q^{j-1} = \frac{1 - q^j}{1 - q} \qquad [j]_q! \equiv [j]_q [j-1]_q \cdots [1]_q.$$

The q-binomial coefficient for integers $k, n \in \mathbb{Z}_+$ $(k \leq n)$ is

$$\begin{bmatrix} n \\ k \end{bmatrix}_q \equiv \frac{[n]_q!}{[k]_q! [n-k]_q!}.$$

Since our purpose is to consider how market size affects the distribution of marriages under μ^{SD}, we make this dependence on M and W explicit. Analogous to the notation used in Section 3, for any prices $p = (p_{\mathcal{M}}, p_{\mathcal{W}})$, we let random variable $K^{\text{SD}}_{p,M,W}$ denote the number of marriages created by μ^{SD} for a random economy at prices p.

THEOREM 3 (Distribution of marriages under μ^{SD}). *For any M, W and prices $p = p_{\mathcal{M}}, p_{\mathcal{W}}$ with incompatibility parameter $q = q(p)$, the number of marriages created under μ^{SD} for a random economy has the following distribution.[13] For $0 \leq k \leq \min\{M, W\}$,*

$$P\big(K^{\text{SD}}_{p,M,W} = k\big) = (1 - q)^k q^{(M-k)(W-k)} \begin{bmatrix} M \\ k \end{bmatrix}_q \begin{bmatrix} W \\ k \end{bmatrix}_q [k]_q!. \tag{1}$$

Its expectation is provided by Kemp (1998):

$$E\big(K^{\text{SD}}_{p,M,W}\big) = \sum_{j=1}^{\min\{M,W\}} \frac{\big[(1 - q^M) \cdots (1 - q^{M-j+1})\big]\big[(1 - q^W) \cdots (1 - q^{W-j+1})\big]}{1 - q^j}. \tag{2}$$

The first two terms in (1) have a straightforward interpretation: the term $(1 - q)^k$ is the probability that k given pairs of agents are compatible, while $q^{(M-k)(W-k)}$ is the probability that all other agents are incompatible. The remaining terms are a probabilistic analog to the number of ways to form k pairs from the sets M, W.

Figure 1 graphs (2) for various levels of incompatibility $q(p)$, fixing $M = 50$ while varying W. The graph illustrates the intuitive fact that when the market is very imbalanced (W far from 50), the platform creates close to the maximum feasible number of marriages ($\min\{W, 50\}$) even at relatively high prices ($q(p)$ close to 1). In relatively balanced markets, alternatively, the platform faces a richer trade-off between price and volume.[14] The same intuition holds true for stable platforms (μ^{DA}).

---

[13]Equation (1) is analyzed by Blomqvist (1952) and Kemp (1998). Ebrahimy and Shimer (2010) use it to describe employment in a stock-flow labor model.

[14]To illustrate, when $M = W = 50$, and $F_{\mathcal{M}}$ and $F_{\mathcal{W}}$ are both $U[0, 1]$, revenue-maximizing prices under μ^{SD} are approximately $p^*_{\mathcal{M}} = p^*_{\mathcal{W}} = 0.718$ ($q(p^*) \approx 0.92$). These prices yield approximately 41.9 expected marriages, leaving 16% of the market unserved. While these numbers are merely illustrative, they demonstrate a nontrivial price–volume trade-off.
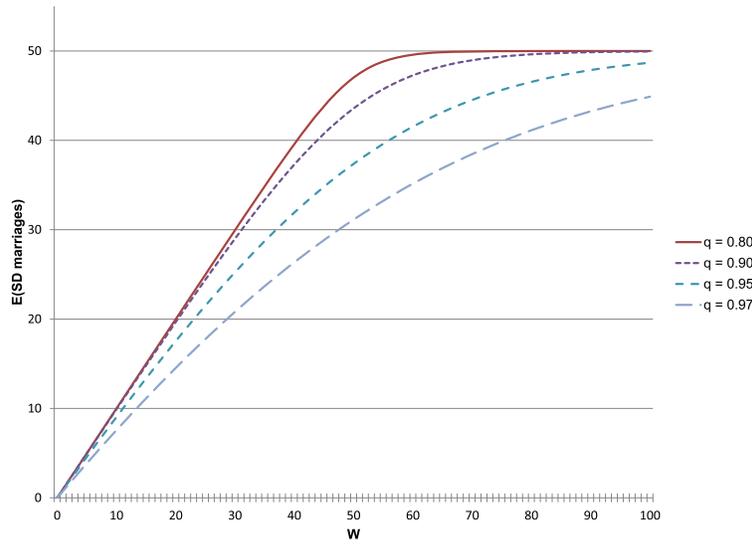
FIGURE 1. The expected number of marriages under serial dictatorship for various $q(p)$, fixing $M = 50$.

We compare the expected number of marriages created under μ^{SD} and μ^{DA}, observing that the former is both an approximation and a bound for the latter. Following the intuition discussed above, it is unsurprising that these values approximate each other in very unbalanced markets, where both matching processes create close to the maximum feasible number of marriages, $\min\{M, W\}$. Therefore, we focus on the ("worst") case of balanced markets ($M = W$), where the expected number of marriages under either process need not be close to $\min\{M, W\}$. Our resulting conclusions easily extend to unbalanced markets.

Figure 2 graphs the percentage by which the expected number of marriages under μ^{SD} falls short of the expected number of marriages under μ^{DA}. To be precise, let $K^{\mathrm{DA}}{}_{p,M,W}$ denote the number of marriages created by μ^{DA} for a random economy at prices p. The figure graphs the percentage $[E(K^{\mathrm{DA}}_{p,M,W}) - E(K^{\mathrm{SD}}_{p,M,W})]/E(K^{\mathrm{DA}}_{p,M,W})$ as a function of (balanced) market size $n = M = W$. Consistent with the related asymptotic results of Arnosti (2016), the values in the graph are positive.

The figure also hints that for any pair of prices (i.e., any $q(p)$), the percentage difference converges to zero. The following theorem implies a slightly stronger conclusion: the expected *number* of *single* agents under μ^{SD} converges to a constant as $n = M = W$ grows large; see Figure 3.[15]

THEOREM 4 (Expected singles under μ^{SD}). *For random, balanced economies of size $n = M = W$, let random variable $S^{\mathrm{SD}}_{p,n}$ denote the number of unmatched men (hence unmatched women) under μ^{SD} at prices p. For any prices p with incompatibility parameter*

---

[15]The distribution of $S^{\mathrm{SD}}_{p,n}$ has thin tails (see (8) in the Appendix). Therefore, Figure 3 approximates $\bar{S}(q)$ by summing the first (sufficiently many) terms of (3).
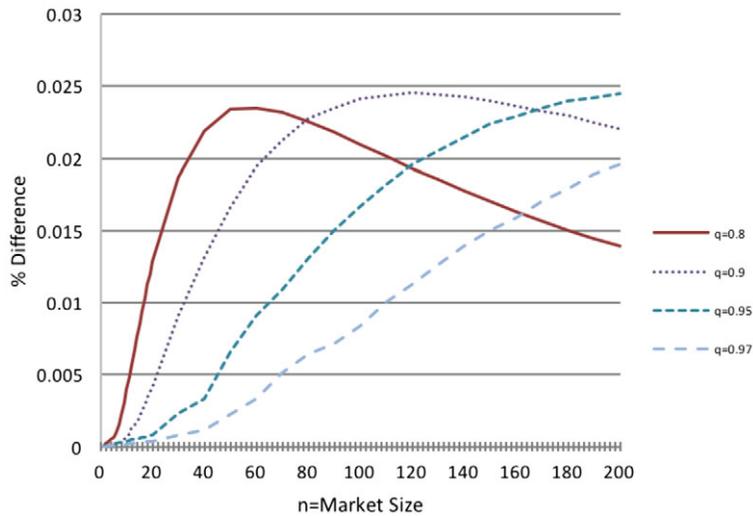
FIGURE 2. The percentage difference in expected marriages between μ^{DA} (simulation) and μ^{SD} (Equation (2)).
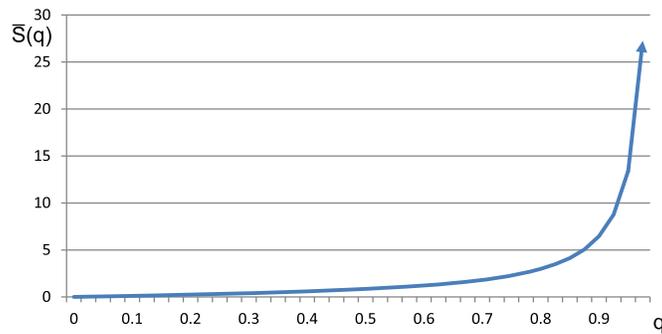

FIGURE 3. The expected number of unmatched men (women) under μ^{SD} in large, balanced markets at prices p is $\bar{S}(q(p))$.

$q = q(p)$, *its asymptotic expectation is*

$$\lim_{M=W\to\infty} E\big(S_{p,n}^{\mathrm{SD}}\big) = \bar{S}(q) \equiv \left[\prod_{i=1}^{\infty}\big(1 - q^i\big)\right]\left[\sum_{s=0}^{\infty} \frac{s \cdot q^{s^2}}{\big((1-q)\cdots(1-q^s)\big)^2}\right]. \qquad (3)$$

The result is significant since it allows us to draw a related conclusion about the number of unmatched agents under μ^{DA}, even in economies of arbitrary size and imbalance. First, begin with a balanced economy of size $n = M = W$ and fix prices p. It is intuitive that $\bar{S}(q)$ is an upper bound for $E(S_{p,n}^{\mathrm{SD}})$; i.e., there are more expected single men/women in the limit than in the finite economy. Second, recall that $E(S_{p,n}^{\mathrm{SD}})$ is an upper bound on the expected number of single men/women under μ^{DA}; i.e., there are fewer marriages under μ^{SD} than under μ^{DA} (Figure 2). Third, under μ^{DA}, it is clear that

the expected number of single men is greater than it would have been in an unbalanced economy of size $n = M' < W'$; i.e., adding agents to one side of the market (say, W) decreases the platform's "shortfall" relative to the maximum feasible number of marriages, $n = \min\{M', W'\}$.

Combining these three inequalities yields our conclusion: the expected number of marriages that a stable platform fails to produce by charging prices $p = (p_{\mathcal{M}}, p_{\mathcal{W}}) > (0, 0)$ is bounded above by $\bar{S}(q)$. Fixing prices, the fraction of potential marriages that the platform fails to create—but that could have been created at lower prices—becomes vanishingly small as the market grows. Not only does a larger market benefit the platform in the obvious way of increasing the number of potential matches, but also does so by increasing the feasible per capita value of matchings yielding higher surplus extraction per agent. The relative "cost of stability" suffered by the platform in smaller markets vanishes as the market grows large.

These conclusions rely on certain assumptions we have made so far. Aside from the standard abstractions we make in our model (e.g., a lack of search frictions, etc.), the most significant assumption has been that of independent preferences (u). We address this in the following section.

5. Preference homogeneity

We reconsider the effect of market imbalance on price allocation when preferences exhibit homogeneity. Recall that under heterogeneous preferences (independently drawn values), market imbalance has no direct effect on price allocation (Theorem 2) and any indirect effect (via its effect on the price level) moves both sides' prices in the same direction (Proposition 2).

Here we show how two forms of preference homogeneity alter these conclusions. First, market imbalance biases price allocation, but in a direction that depends on the which type of homogeneity is present (Proposition 3). Second, any change in market imbalance moves the two sides' prices in opposite directions (Proposition 4).

We consider the two forms of homogeneity separately. In the first—*within-side homogeneity*—all agents on one side of the market agree on the desirability of any given agent on the other side. Such "vertical" preferences have been assumed in various works on exclusive matching platforms discussed in Section 1.2. In the second—*partner homogeneity*—any given agent values all potential partners equally. This assumption, appearing, for example, in the two-sided markets literature, reflects undifferentiated partners while allowing heterogeneity in outside options.

Throughout this section, we assume that the profile of agents' values u is drawn from a joint distribution that satisfies one of the following two sets of assumptions.

DEFINITION 4. Preferences exhibit *within-side homogeneity* when the following conditions occur.

- For each $w \in W$, the men have a common value $U_{\mathcal{M}}(w)$ drawn from $F_{\mathcal{M}}$, so for all $m \in M$, $u_m(w) = U_{\mathcal{M}}(w)$.

- For each $m \in M$, the women have a common value $U_{\mathcal{W}}(m)$ drawn from $F_{\mathcal{W}}$, so for all $w \in W$, $u_w(m) = U_{\mathcal{W}}(m)$.

- These $(W + M)$ different values are drawn independently.

Preferences exhibit *partner homogeneity* when the following conditions occur.

- Each $m \in M$ has a participation value $U_{\mathcal{M}}(m)$ drawn from $F_{\mathcal{M}}$, so for all $w \in W$, $u_m(w) = U_{\mathcal{M}}(m)$.

- Each $w \in W$ has a participation value $U_{\mathcal{W}}(w)$ drawn from $F_{\mathcal{W}}$, so for all $m \in M$, $u_w(m) = U_{\mathcal{W}}(w)$.

- These $(M + W)$ different values are drawn independently.

Under either form of preference homogeneity, the space of plausible matching processes becomes less rich. Indeed much work under within-side homogeneity assumes or derives a stable (assortative) matching outcome. Our observations apply more generally to any matching processes that satisfy Condition 3, since they yield the same number of marriages as μ^{DA}.

FACT. *Suppose preferences exhibit either form of homogeneity. Let μ be individually rational and weakly unimprovable. For any prices p, the expected number of marriages under μ and μ^{DA} is equal.*

This fact is proven simply. Under within-side homogeneity, the women commonly find any given man "acceptable" with probability $1 - F_{\mathcal{W}}(p_{\mathcal{W}})$. The number of acceptable men is thus a (binomial) random variable, $k_{\mathcal{M}}$. Similarly, the number of acceptable women is $k_{\mathcal{W}}$. Clearly μ^{DA} creates $k = \min\{k_{\mathcal{M}}, k_{\mathcal{W}}\}$ marriages (assortatively). An individually rational μ creates no more than k marriages; a weakly unimprovable μ creates no fewer.

Analogously, under partner homogeneity, there are (binomially distributed) $k'_{\mathcal{M}}$ men and $k'_{\mathcal{W}}$ women "willing to participate." An individually rational, unimprovable μ creates $k' = \min\{k'_{\mathcal{M}}, k'_{\mathcal{W}}\}$ marriages.

In either case, expected revenue maximization involves the (intractable) expected minimum of two binomial variables. Since our goal is to qualitatively contrast market imbalance effects under our two forms of correlation, we simplify the discussion by examining a large (continuum) market.

5.1 *Price allocation: Large markets with homogeneity*

We consider the model of Section 2 but with a continuum of agents: a mass $\tilde{M}$ of men and a mass $\tilde{W}$ of women. Our definitions extend to this setting in a straightforward way, so we omit their reformalization for brevity.

Under within-side homogeneity, analogous to $k_{\mathcal{M}}$ and $k_{\mathcal{W}}$ above, there are (now deterministic) masses of "acceptable" men $\kappa_{\mathcal{M}} = (1 - F_{\mathcal{W}}(p_{\mathcal{W}})) * \tilde{M}$ and women $\kappa_{\mathcal{W}} = (1 -$
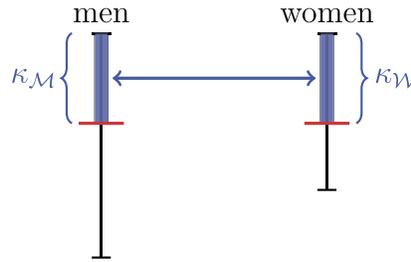
FIGURE 4. Within-side homogeneity leads to a relatively higher price for the short side of the market: $\tilde{M} > \tilde{W}$ leads to $F_{\mathcal{W}}(p_{\mathcal{W}}) > F_{\mathcal{M}}(p_{\mathcal{M}})$.

$F_{\mathcal{M}}(p_{\mathcal{M}})) * \tilde{W}$, yielding a mass of marriages $\kappa = \min\{\kappa_{\mathcal{M}}, \kappa_{\mathcal{W}}\}$. Revenue-maximizing prices clearly yield $\kappa_{\mathcal{M}} = \kappa_{\mathcal{W}}$, implying

$$\frac{1 - F_{\mathcal{W}}(p_{\mathcal{W}})}{1 - F_{\mathcal{M}}(p_{\mathcal{M}})} = \frac{\tilde{W}}{\tilde{M}}. \tag{4}$$

Therefore (see Figure 4), the short side of the market is charged a relatively higher price than the long side in the sense that if, say, $\tilde{W} < \tilde{M}$, then $F_{\mathcal{W}}(p_{\mathcal{W}}) > F_{\mathcal{M}}(p_{\mathcal{M}})$. When $F_{\mathcal{M}} = F_{\mathcal{W}}$, this additionally means $p_{\mathcal{W}} > p_{\mathcal{M}}$.

Partner homogeneity inverts this relationship: The platform equates the masses of men $\kappa'_{\mathcal{M}} = (1 - F_{\mathcal{W}}(p_{\mathcal{W}})) * \tilde{W}$ and women $\kappa'_{\mathcal{W}} = (1 - F_{\mathcal{M}}(p_{\mathcal{M}})) * \tilde{M}$ who are "willing to participate." By inverting the right hand side of (4), the price bias now relatively *favors* the short side. We summarize as follows.

PROPOSITION 3 (Market imbalance and preference homogeneity). *Consider an imbalanced market where (without loss of generality) $\tilde{M} > \tilde{W}$ and suppose the matching process is individually rational and weakly unimprovable.*

- *Under within-side homogeneity, the revenue-maximizing platform charges a relatively higher price to the short side: $F_{\mathcal{W}}(p_{\mathcal{W}}) > F_{\mathcal{M}}(p_{\mathcal{M}})$.*

- *Under partner homogeneity, the revenue-maximizing platform charges a relatively higher price to the long side: $F_{\mathcal{M}}(p_{\mathcal{M}}) > F_{\mathcal{W}}(p_{\mathcal{W}})$.*

The proof follows from (4).

Last, we show that with preference homogeneity, a change in market imbalance moves the two sides' prices in opposite directions. This contrasts with the independent values case where prices co-vary (Proposition 2).

PROPOSITION 4 (Price negative co-movement). *Suppose $F_{\mathcal{M}}$ and $F_{\mathcal{W}}$ have strictly increasing hazard rates. An increase in the proportion of men, $\tilde{M}/\tilde{W}$, changes revenue-maximizing prices $(p^*_{\mathcal{M}}, p^*_{\mathcal{W}})$ as follows.*

- *Within-side homogeneity: $p^*_{\mathcal{M}}$ weakly decreases and $p^*_{\mathcal{W}}$ weakly increases.*

- *Partner homogeneity: $p^*_{\mathcal{M}}$ weakly increases and $p^*_{\mathcal{W}}$ weakly decreases.*

An intuitive illustration of this result is so-called surge pricing on ride-sharing platforms that exhibit partner homogeneity (agents are ex ante indifferent over partners). A relative shortage of drivers leads to higher rider prices and lower driver prices (higher wage). Interestingly, the introduction of heterogeneous preferences over potential partners to such environments would make surge pricing less desirable (Proposition 2).

For within-side homogeneity, an increase in $\tilde{M}/\tilde{W}$ gives each woman a more valuable partner. As the platform extracts this value by raising $p_{\mathcal{W}}$, marriages become more profitable. The platform thus lowers $p_{\mathcal{M}}$ to generate more of them. This argument is reminiscent of the "see-saw" effect in Rochet and Tirole (2006): a price increase on one side incentivizes a price decrease on the other. Unlike in their model, the argument holds here only under an *interaction* between preference homogeneity and market imbalance.

6. Conclusion

We have established a qualitative distinction between pricing decisions on exclusive and nonexclusive matching platforms. With exclusivity, an interaction between preference homogeneity and market imbalance impacts optimal price allocation: transaction prices become biased toward either the short or long side of the market depending on the form of homogeneity. This bias disappears when preferences are fully heterogeneous. An amplification in market imbalance either moves the two sides' prices in the same direction (heterogeneous preferences) or opposite direction (homogeneous preferences).

The reasoning behind this is that the two cases differ in their price allocation objective. For ordinal, monotonic matching processes, this objective under heterogeneous preferences is to minimize the incompatibility rate for any arbitrary pair of agents (Lemma 1). Under homogeneous preferences the objective is to minimize incompatibility between "marginal" participants on the platform (Equation (4)), which depends on market sizes.

These results highlight an empirical requirement on exclusive matching platforms that is absent in the case of nonexclusive matching: knowledge of the presence and structure of preference homogeneity and, unless preferences are sufficiently heterogeneous, an assessment of the imbalance between the sizes of the two sides of the market.

Appendix: Omitted proofs and examples

The following example illustrates that the conclusions of Section 3 may fail for non-ordinal matching processes.

Example 1 (Asymmetric pricing for non-ordinal μ). Let $M = \{m\}$ and $W = \{1, \ldots, 10\}$, with each value $u_i(j)$ drawn uniformly from $[0, 1]$. Let μ be such that, for any prices p and realized values u, m marries his favorite woman, $\arg\max u_m(w)$, if she is p-compatible with him, and otherwise remains single. The outcome is sensitive to m's preferences over incompatible women, so μ is not ordinal. It is easy to see that (i) at

prices $p_{\mathcal{M}} = 0$ and $p_{\mathcal{W}} = 0.5$, a marriage is created with probability 0.5; (ii) at prices $p'_{\mathcal{M}} = 0.5$ and $p'_{\mathcal{W}} = 0$ this probability is $(1 - 0.5^{10})$. Even though $q(p) = q(p')$, the expected number of marriages is greater at prices p'. ◇

Proposition 1 and Proposition 2 are proven with the following result.

LEMMA 2. *Suppose the matching process is ordinal and monotonic, and that $F_{\mathcal{M}}$ and $F_{\mathcal{W}}$ have strictly increasing hazard rates. For any $0 < p_T < 2$, there is a unique pair of prices, $(p^*_{\mathcal{M}}, p^*_{\mathcal{W}})$, that maximize expected revenue subject to the constraint $p_{\mathcal{M}} + p_{\mathcal{W}} = p_T$. These prices minimize the absolute difference in the two sides' hazard rates:*

$$(p^*_{\mathcal{M}}, p^*_{\mathcal{W}}) = \underset{p_{\mathcal{M}}, p_{\mathcal{W}} \,:\, p_{\mathcal{M}} + p_{\mathcal{W}} = p_T}{\arg\min} \big| h_{\mathcal{M}}(p_{\mathcal{M}}) - h_{\mathcal{W}}(p_{\mathcal{W}}) \big|. \tag{5}$$

PROOF. By Theorem 2 it suffices to show that there is a unique pair of prices that minimize $q(p_{\mathcal{M}}, p_{\mathcal{W}})$ subject to the constraint $p_{\mathcal{M}} + p_{\mathcal{W}} = p_T$, and that these prices satisfy (5). Recall our restriction to prices $p_{\mathcal{M}}, p_{\mathcal{W}} \in [0, 1]$. Note that this restriction is without loss of generality here since, for $0 < p_T < 2$, q-minimizing prices must satisfy $p_{\mathcal{M}}, p_{\mathcal{W}} \in [0, 1)$. First, $p_{\mathcal{M}} \geq 1$ or $p_{\mathcal{W}} \geq 1$ implies $q() = 1$ while prices $(p_T/2, p_T/2)$ guarantee $q() < 1$. Second, if say $p_{\mathcal{M}} < 0$, then a small increase in $p_{\mathcal{M}}$ and decrease in $p_{\mathcal{W}} = p_T - p_{\mathcal{M}} > 0$ strictly decreases $q()$.

Under the constraint $p_{\mathcal{M}} + p_{\mathcal{W}} = p_T$, the requirement $p_{\mathcal{M}}, p_{\mathcal{W}} \in [0, 1]$ can be written as $\max\{0, p_T - 1\} \leq p_{\mathcal{M}} \leq \min\{1, p_T\}$. Therefore, defining $\tilde{q}(p_{\mathcal{M}}) \equiv q(p_{\mathcal{M}}, p_T - p_{\mathcal{M}})$, the minimization problem can be written as

$$\min_{p_{\mathcal{M}}} \tilde{q}(p_{\mathcal{M}}) \quad \text{s.t. } \max\{0, p_T - 1\} \leq p_{\mathcal{M}} \leq \min\{1, p_T\}. \tag{6}$$

Differentiating $\tilde{q}$ on this range, we have

$$\frac{d\tilde{q}(p_{\mathcal{M}})}{dp_{\mathcal{M}}} = \big(1 - F_{\mathcal{W}}(p_T - p_{\mathcal{M}})\big) f_{\mathcal{M}}(p_{\mathcal{M}}) - \big(1 - F_{\mathcal{M}}(p_{\mathcal{M}})\big) f_{\mathcal{W}}(p_T - p_{\mathcal{M}}),$$

which is continuous in $p_{\mathcal{M}}$. For interior values of $p_{\mathcal{M}}$, i.e., $\max\{0, p_T - 1\} < p_{\mathcal{M}} < \min\{1, p_T\}$, this can be written as

$$\frac{d\tilde{q}(p_{\mathcal{M}})}{dp_{\mathcal{M}}} = \big(1 - F_{\mathcal{W}}(p_T - p_{\mathcal{M}})\big)\big(1 - F_{\mathcal{M}}(p_{\mathcal{M}})\big)\big(h_{\mathcal{M}}(p_{\mathcal{M}}) - h_{\mathcal{W}}(p_T - p_{\mathcal{M}})\big),$$

where division by zero is avoided for $p_{\mathcal{M}} < \min\{1, p_T\}$. Furthermore, this inequality along with $\max\{0, p_T - 1\} < p_{\mathcal{M}}$ also implies

$$\big(1 - F_{\mathcal{W}}(p_T - p_{\mathcal{M}})\big)\big(1 - F_{\mathcal{M}}(p_{\mathcal{M}})\big) > 0,$$

meaning that for interior values of $p_{\mathcal{M}}$, $d\tilde{q}(p_{\mathcal{M}})/dp_{\mathcal{M}}$ has the same sign as the difference

$$\big(h_{\mathcal{M}}(p_{\mathcal{M}}) - h_{\mathcal{W}}(p_T - p_{\mathcal{M}})\big). \tag{7}$$

By the monotone hazard rate assumption, (7) is strictly increasing in $p_{\mathcal{M}}$. Hence on the range $\max\{0, p_T - 1\} \leq p_{\mathcal{M}} \leq \min\{1, p_T\}$, the sign of $d\tilde{q}(\cdot)/dp_{\mathcal{M}}$ is either (i) always negative, (ii) always positive, or (iii) crosses zero from below at exactly one price $p_{\mathcal{M}}^*$. In these three respective cases, $\tilde{q}(\,)$ is minimized at (i) $p_{\mathcal{M}}^* = \min\{1, p_T\}$, (ii) $p_{\mathcal{M}}^* = \max\{0, p_T - 1\}$, or (iii) where $h_{\mathcal{M}}(p_{\mathcal{M}}^*) = h_{\mathcal{W}}(p_T - p_{\mathcal{M}}^*)$. In each case, this is the price that minimizes (5), proving the lemma.   □

PROOF OF PROPOSITION 1. The result follows immediately from Lemma 2 and observing that the solution to (5) is given by $p_{\mathcal{M}}^*(p_T) = p_T/2$ when $F_{\mathcal{M}} = F_{\mathcal{W}}$.   □

Without the hazard rate condition, revenue-maximizing prices may be unequal even if $F_{\mathcal{M}} = F_{\mathcal{W}}$, as illustrated in the following discretized example.

EXAMPLE 2 (Optimal, unequal prices). Consider one man and one woman ($M = W = 1$), and the matching process that creates a marriage whenever the agents are p-compatible. The value that each agent assigns to the potential mate is (independently) either 0.1 (probability π) or 0.9 (probability $1 - \pi$). One can restrict attention to prices $p_{\mathcal{M}}, p_{\mathcal{W}} \in \{0.1, 0.9\}$ and check by inspection that the following price pairs maximize expected revenue:

$$(p_{\mathcal{M}}^*, p_{\mathcal{W}}^*) = (0.9, 0.9) \quad \text{when } \pi \leq 4/9$$
$$(p_{\mathcal{M}}^*, p_{\mathcal{W}}^*) \in \{(0.1, 0.9), (0.9, 0.1)\} \quad \text{when } 4/9 \leq \pi \leq 4/5$$
$$(p_{\mathcal{M}}^*, p_{\mathcal{W}}^*) = (0.1, 0.1) \quad \text{when } 4/5 \leq \pi.$$

In the case $4/9 < \pi < 4/5$, it is strictly optimal to charge unequal prices. Nevertheless, the *set* of optimal price lists is symmetric (see Theorem 1).   ◊

PROOF OF PROPOSITION 2. For any $0 < p_T < 2$, Lemma 2 states that there are unique revenue-maximizing prices, $(p_{\mathcal{M}}^*(p_T), p_{\mathcal{W}}^*(p_T))$. We show that $p_{\mathcal{M}}^*(p_T)$ is nondecreasing in p_T. An identical argument applies to $p_{\mathcal{W}}^*(p_T)$.

Fix $0 < p_T' < p_T'' < 2$ and denote optimal price allocations

$$p_{\mathcal{M}}' = p_{\mathcal{M}}^*(p_T') \qquad p_{\mathcal{W}}' = p_T' - p_{\mathcal{M}}^*(p_T') \qquad p_{\mathcal{M}}'' = p_{\mathcal{M}}^*(p_T'') \qquad p_{\mathcal{W}}'' = p_T'' - p_{\mathcal{M}}^*(p_T'').$$

Suppose, to the contrary, that $p_{\mathcal{M}}'' = p_{\mathcal{M}}' - \delta$ for some $\delta > 0$. With the constraints in (6), this implies

$$\max\{0, p_T'' - 1\} \leq p_{\mathcal{M}}'' < p_{\mathcal{M}}' \leq \min\{1, p_T'\}$$
$$\max\{0, p_T' - 1\} \leq p_{\mathcal{W}}' < p_{\mathcal{W}}' + \delta < p_{\mathcal{W}}'' \leq \min\{1, p_T''\}.$$

Next observe that

$$h_{\mathcal{M}}(p_{\mathcal{M}}'') < h_{\mathcal{M}}(p_{\mathcal{M}}') \leq h_{\mathcal{W}}(p_{\mathcal{W}}') < h_{\mathcal{W}}(p_{\mathcal{W}}' + \delta) < h_{\mathcal{W}}(p_{\mathcal{W}}'').$$

The strict inequalities follow immediately from the hazard rate assumption. To derive the weak inequality, observe that if $h_{\mathcal{M}}(p'_{\mathcal{M}}) > h_{\mathcal{W}}(p'_{\mathcal{W}})$, then for small $\epsilon > 0$, $p'_{\mathcal{M}} - \epsilon > 0$ and $p'_{\mathcal{W}} + \epsilon < 1$ would strictly reduce the absolute difference in hazard rates, i.e.,

$$\left| h_{\mathcal{M}}(p'_{\mathcal{M}} - \epsilon) - h_{\mathcal{W}}(p'_{\mathcal{W}} + \epsilon) \right| < \left| h_{\mathcal{M}}(p'_{\mathcal{M}}) - h_{\mathcal{W}}(p'_{\mathcal{W}}) \right|,$$

in contradiction to (5).

By definition, $(p''_{\mathcal{M}}, p''_{\mathcal{W}})$ minimizes (5) with respect to p''_T. However, $h_{\mathcal{M}}(p''_{\mathcal{M}}) < h_{\mathcal{W}}(p''_{\mathcal{W}})$ implies that for small $\epsilon > 0$, $p''_{\mathcal{M}} + \epsilon < 1$ and $p''_{\mathcal{W}} - \epsilon > 0$ strictly reduces the absolute difference in hazard rates, i.e.,

$$\left| h_{\mathcal{M}}(p''_{\mathcal{M}} + \epsilon) - h_{\mathcal{W}}(p''_{\mathcal{W}} - \epsilon) \right| < \left| h_{\mathcal{M}}(p''_{\mathcal{M}}) - h_{\mathcal{W}}(p''_{\mathcal{W}}) \right|,$$

which contradicts (5). $\qquad\square$

PROOF OF THEOREM 3. Fix $p = (p_{\mathcal{M}}, p_{\mathcal{W}})$ with incompatibility parameter $q = q(p)$. We prove (1) for arbitrary M by induction on W. Since p is fixed, we simplify notation by writing $K_{M,W}$ to denote $K^{\mathrm{SD}}_{p,M,W}$.

Fix M. Equation (1) clearly holds when $W = 1$: $P(K_{M,1} = 0) = q^M$ (the woman is incompatible with each man) and $P(K_{M,1} = 1) = 1 - q^M$. For some arbitrary W, suppose (1) accurately describes the distribution of $K_{M,W-1}$. We show that it accurately describes the distribution of $K_{M,W}$.

Observe that the SD algorithm (Definition 2) creates at most one (permanent) marriage in each step $1 \le t \le W$. Also, the total number of marriages created through step $W - 1$ has the same distribution as $K_{M,W-1}$. In words, running SD on a random economy of size $(M, W - 1)$ is equivalent to running the first $W - 1$ steps of SD on a random economy of size (M, W).

Furthermore, for SD to yield a total of k marriages after the final step $t = W$, there must have been either k or $k - 1$ marriages created through step $t = W - 1$. Consider the (conditional) probability of obtaining k marriages given either of these two scenarios.

Scenario 1: k marriages are created through step $t = W - 1$. Entering step $t = W$, there are $M - k$ men yet unmatched. Woman W is incompatible with each of them with probability q^{M-k}. That is, with probability q^{M-k}, step $t = W$ adds no additional marriage to the existing k.

Scenario 2: $k - 1$ marriages are created through step $t = W - 1$. Entering step $t = W$, there are $M - k + 1$ men currently unmatched. Woman W is p-compatible with at least one with probability $1 - q^{M-k+1}$. That is, with probability $1 - q^{M-k+1}$, step $t = W$ adds a kth marriage to the existing $k - 1$.

The total probability of SD yielding k marriages is, therefore,

$$P(K_{M,W} = k) = P(K_{M,W-1} = k - 1)\left(1 - q^{M-k+1}\right) + P(K_{M,W-1} = k)q^{M-k}.$$

Using (1) to substitute for the distribution of $K_{M,W-1}$, this equals

$$(1-q)^{k-1}q^{(M-k+1)(W-k)}\begin{bmatrix}M\\k-1\end{bmatrix}_q\begin{bmatrix}W-1\\k-1\end{bmatrix}_q[k-1]_q!(1-q^{M-k+1})$$

$$+(1-q)^k q^{(M-k)(W-k-1)}\begin{bmatrix}M\\k\end{bmatrix}_q\begin{bmatrix}W-1\\k\end{bmatrix}_q[k]_q!(q^{M-k})$$

$$=\left(\frac{q^{W-k}}{[k]_q(1-q)}\right)(1-q)^k q^{(M-k)(W-k)}\begin{bmatrix}M\\k-1\end{bmatrix}_q\begin{bmatrix}W-1\\k-1\end{bmatrix}_q[k]_q!(1-q^{M-k+1})$$

$$+(1-q)^k q^{(M-k)(W-k)}\begin{bmatrix}M\\k\end{bmatrix}_q\begin{bmatrix}W-1\\k\end{bmatrix}_q[k]_q!$$

$$=\left(\frac{q^{W-k}}{[k]_q(1-q)}\right)(1-q)^k q^{(M-k)(W-k)}\begin{bmatrix}M\\k\end{bmatrix}_q\frac{[k]_q}{[M-k+1]_q}\begin{bmatrix}W\\k\end{bmatrix}_q\frac{[k]_q}{[W]_q}[k]_q!(1-q^{M-k+1})$$

$$+(1-q)^k q^{(M-k)(W-k)}\begin{bmatrix}M\\k\end{bmatrix}_q\begin{bmatrix}W\\k\end{bmatrix}_q\frac{[W-k]_q}{[W]_q}[k]_q!$$

$$=q^{W-k}(1-q)^k q^{(M-k)(W-k)}\begin{bmatrix}M\\k\end{bmatrix}_q\begin{bmatrix}W\\k\end{bmatrix}_q\frac{[k]_q}{[W]_q}[k]_q!$$

$$+(1-q)^k q^{(M-k)(W-k)}\begin{bmatrix}M\\k\end{bmatrix}_q\begin{bmatrix}W\\k\end{bmatrix}_q\frac{[W-k]_q}{[W]_q}[k]_q!$$

$$=(1-q)^k q^{(M-k)(W-k)}\begin{bmatrix}M\\k\end{bmatrix}_q\begin{bmatrix}W\\k\end{bmatrix}_q[k]_q!\left(\frac{q^{W-k}[k]_q+[W-k]_q}{[W]_q}\right)$$

$$=(1-q)^k q^{(M-k)(W-k)}\begin{bmatrix}M\\k\end{bmatrix}_q\begin{bmatrix}W\\k\end{bmatrix}_q[k]_q!,$$

proving (1) for $K_{M,W}$. Equation (2) is proven by Kemp (1998). $\qquad\square$

PROOF OF THEOREM 4. Fix prices p with $q=q(p)$. In a balanced market of size $n=M=W$, the probability that μ^{SD} yields a *perfect matching*, i.e., of having zero unmatched agents, is given by (1):

$$P\big(K_{p,n,n}^{\mathrm{SD}}=n\big)=(1-q)^n q^{(n-n)(n-n)}\begin{bmatrix}n\\n\end{bmatrix}_q\begin{bmatrix}n\\n\end{bmatrix}_q[n]_q!$$

$$=\big(1-q^n\big)\cdots\big(1-q^1\big).$$

As n goes to infinity, $P(K_{p,n,n}^{\mathrm{SD}}=n)$ converges to $\phi(q)\equiv\prod_{i=1}^{\infty}(1-q^i)$. It can be shown that $0\le q<1$ implies $\phi(q)>0$. Therefore, $P(K_{p,n,n}^{\mathrm{SD}}=n)$ is bounded away from zero across all market sizes n.

In unbalanced markets of arbitrary sizes M, W, the probability of exactly k marriages, and hence $g=M-k$ single men and $h=W-k$ single women, can be written in

terms of g and h using (1):

$$P\big(K^{\text{SD}}_{p,M,W}=k\big)=P\big(K^{\text{SD}}_{p,k+g,k+h}=k\big)=(1-q)^k q^{gh}\begin{bmatrix}k+g\\k\end{bmatrix}_q\begin{bmatrix}k+h\\k\end{bmatrix}_q[k]_q!.$$

Letting the market grow large by letting $k\to\infty$, the probability that there are g single men (and hence h single women) converges to

$$\lim_{k\to\infty}P\big(K^{\text{SD}}_{p,k+g,k+h}=k\big)$$

$$=\lim_{k\to\infty}(1-q)^k q^{gh}\begin{bmatrix}k+g\\k\end{bmatrix}_q\begin{bmatrix}k+h\\k\end{bmatrix}_q[k]_q!$$

$$=\lim_{k\to\infty}(1-q)^k q^{gh}\left(\frac{\big(1-q^{k+1}\big)\cdots\big(1-q^{k+g}\big)}{(1-q)\cdots\big(1-q^g\big)}\right)\left(\frac{\big(1-q^{k+1}\big)\cdots\big(1-q^{k+h}\big)}{(1-q)\cdots\big(1-q^h\big)}\right)[k]_q!$$

$$=\lim_{k\to\infty}\big(1-q^k\big)\cdots(1-q)q^{gh}\left(\frac{\big(1-q^{k+1}\big)\cdots\big(1-q^{k+g}\big)}{(1-q)\cdots\big(1-q^g\big)}\right)\left(\frac{\big(1-q^{k+1}\big)\cdots\big(1-q^{k+h}\big)}{(1-q)\cdots\big(1-q^h\big)}\right)$$

$$=\phi(q)q^{gh}\frac{1}{(1-q)\cdots\big(1-q^g\big)\cdot(1-q)\cdots\big(1-q^h\big)}. \tag{8}$$

Hence in the case of asymptotically large balanced markets ($g=h$), the expected number of single men (and hence women) converges to (3):

$$\bar{S}(q)=\phi(q)\sum_{s=0}^{\infty}s\cdot q^{ss}\frac{1}{(1-q)\cdots\big(1-q^s\big)\cdot(1-q)\cdots\big(1-q^s\big)}. \qquad\square$$

PROOF OF PROPOSITION 4. We prove that under partner homogeneity, the revenue-maximizing price charged to the men is weakly increasing in $\eta\equiv\tilde{M}/\tilde{W}$, the relative proportion of men. By relabeling the sides of the market, the same proof implies that the women's revenue-maximizing price weakly decreases in η.

For $i=\mathcal{M},\mathcal{W}$ and $x\in[0,1]$, define $\tilde{p}_i(x)\equiv F_i^{-1}(1-x)$ to be the price at which x is the proportion of agents on side i who are willing to match. Our assumptions on $F_{\mathcal{M}}$ and $F_{\mathcal{W}}$ imply that $\tilde{p}_{\mathcal{M}}()$ and $\tilde{p}_{\mathcal{W}}()$ are strictly decreasing and continuously differentiable.

Recall that the platform sets prices in a way that equates the two masses of agents willing to match, which we write as $\kappa_{\mathcal{M}}=\kappa_{\mathcal{W}}=\kappa$, a là Figure 4. Revenue maximization can be written as the optimal choice of κ:

$$\max_{\kappa\le\min\{\tilde{M},\tilde{W}\}}\kappa\cdot\big[\tilde{p}_{\mathcal{M}}(\kappa/\tilde{M})+\tilde{p}_{\mathcal{W}}(\kappa/\tilde{W})\big].$$

Equivalently, the objective can be written in terms of choosing the *proportion* of men willing to match, which we denote $\tilde{\kappa}\equiv\kappa/\tilde{M}$:

$$\max_{\tilde{\kappa}\le\min\{1,1/\eta\}}\tilde{M}\tilde{\kappa}\cdot\big[\tilde{p}_{\mathcal{M}}(\tilde{\kappa})+\tilde{p}_{\mathcal{W}}(\eta\tilde{\kappa})\big].$$

The first-order condition necessary for an interior optimum $0 < \tilde{\kappa} < \min\{1, 1/\eta\}$ is

$$\tilde{p}_{\mathcal{M}}(\tilde{\kappa}) + \tilde{p}_{\mathcal{W}}(\eta\tilde{\kappa}) + \tilde{\kappa} \cdot \tilde{p}'_{\mathcal{M}}(\tilde{\kappa}) + \eta\tilde{\kappa} \cdot \tilde{p}'_{\mathcal{W}}(\eta\tilde{\kappa}) = 0. \tag{9}$$

By definition, $x = 1 - F_i(\tilde{p}_i(x))$. By the inverse function theorem $\tilde{p}'_i(x) = -1/f_i(\tilde{p}_i(x))$. Substituting these terms and denoting the hazard rates $h_i(x) = f_i(x)/(1 - F_i(x))$, the first-order condition (9) is

$$G(\tilde{\kappa}, \eta) \equiv \tilde{p}_{\mathcal{M}}(\tilde{\kappa}) + \tilde{p}_{\mathcal{W}}(\eta\tilde{\kappa}) - \frac{1}{h_{\mathcal{M}}(\tilde{p}_{\mathcal{M}}(\tilde{\kappa}))} - \frac{1}{h_{\mathcal{W}}(\tilde{p}_{\mathcal{W}}(\eta\tilde{\kappa}))} = 0.$$

Since each $\tilde{p}_i()$ is strictly decreasing in $\tilde{\kappa}$, the hazard rate assumption implies that $-1/h_i(\tilde{p}_i(\cdot))$ is also strictly decreasing in $\tilde{\kappa}$. Therefore, $G()$ is strictly decreasing in $\tilde{\kappa}$ and the second-order conditions are also satisfied.

Observe that $G(\tilde{\kappa}, \eta)$ is continuous in $\tilde{\kappa}$ and strictly positive when evaluated at $\tilde{\kappa} = 0$. Therefore, there are two possibilities: either there is a unique interior optimizer $0 < \tilde{\kappa} < \min\{1, 1/\eta\}$ that satisfies (9) or $G(\tilde{\kappa}, \eta) > 0$ for all interior $\tilde{\kappa}$, the constraint binds, and the unique optimizer is $\tilde{\kappa} = \min\{1, 1/\eta\}$.

Regardless of the case, let $\tilde{\kappa}^*(\eta)$ denote the unique optimizer as a function of η. Consider an increase in the ratio of men to women from some η' to some $\eta'' > \eta'$. As η increases, $G(\tilde{\kappa}, \eta)$ decreases and the constraint $\tilde{\kappa} \leq \min\{1, 1/\eta\}$ becomes weakly tighter. Therefore, this must lead to a weakly lower optimal proportion of matched men: $\tilde{\kappa}^*(\eta'') \leq \tilde{\kappa}^*(\eta')$, whether the constraint binds or not. This implies a weak increase in the optimal price charged to the men:

$$\tilde{p}_{\mathcal{M}}(\tilde{\kappa}^*(\eta'')) \geq \tilde{p}_{\mathcal{M}}(\tilde{\kappa}^*(\eta')).$$

By relabeling the sides of the market, the same conclusion can be drawn for the optimal price charged to the women's side.

The proof for *within-side homogeneity* is analogous, but relates the proportion of acceptable men $\tilde{\kappa}$ to the women's price rather than the men's. Thus the opposite conclusion is reached, proving the proposition. □

References

Adachi, Hiroyuki (2003), "A search model of two-sided matching under nontransferable utility." *Journal of Economic Theory*, 113, 182–198. [800]

Alcalde, José (1996), "Implementation of stable solutions to marriage problems." *Journal of Economic Theory*, 69, 240–254. [804]

Ambrus, Attila and Rossella Argenziano (2009), "Asymmetric networks in two-sided markets." *American Economic Journal: Microeconomics*, 1, 17–52. [803]

Armstrong, Mark (2006), "Competition in two-sided markets." *RAND Journal of Economics*, 37, 668–691. [802]

Arnosti, Nick (2016), "Centralized clearinghouse design: A quantity-quality tradeoff." Unpublished paper, SSRN 2571527. [801, 809, 811]

Ashlagi, Itai, Yash Kanoria, and Jacob Leshno (2017), "Unbalanced random matching markets: The stark effect of competition." *Journal of Political Economy*, 125, 69–98. [799, 801, 802]

Biglaiser, Gary (1993), "Middlemen as experts." *RAND Journal of Economics*, 24, 212–223. [803]

Bloch, Francis and Harl Ryder (2000), "Two-sided search, marriages, and matchmakers." *International Economic Review*, 41, 93–115. [802]

Blomqvist, Nils (1952), "On an exhaustion process." *Scandinavian Actuarial Journal*, 35, 201–210. [810]

Caillaud, Bernard and Bruno Jullien (2003), "Chicken & egg: Competition among inter-mediation service providers." *RAND Journal of Economics*, 34, 309–328. [803]

Damiano, Ettore and Li Hao (2008), "Competing matchmaking." *Journal of the European Economic Association*, 6, 789–818. [802]

Damiano, Ettore and Hao Li (2007), "Price discrimination and efficient matching." *Economic Theory*, 30, 243–263. [802]

Ebrahimy, Ehsan and Robert Shimer (2010), "Stock-flow matching." *Journal of Economic Theory*, 145, 1325–1353. [810]

Fershtman, Daniel and Alessandro Pavan (2020), "Matching auctions." Unpublished paper, Department of Economics, Northwestern University. [803]

Gale, David and S. Shapley Lloyd (1962), "College admissions and the stability of marriage." *American Mathematical Monthly*, 69, 9–15. [800, 808]

Gale, David and Marilda Sotomayor (1985), "Ms. Machiavelli and the stable matching problem." *American Mathematical Monthly*, 92, 261–268. [805, 809]

Gomes, Renato and Alessandro Pavan (2016), "Many-to-many matching and price discrimination." *Theoretical Economics*, 11, 1005–1052. [803]

Halaburda, Hanna, Jan Piskorski, and Pinar Yildirim (2018), "Competing by restricting choice: The case of matching platforms." *Management Science*, 64, 3574–3594. [803]

Hitch, Gunter J., Ali Hortacsu, and Dan Ariely (2010), "Matching and sorting in online dating." *American Economic Review*, 100, 130–163. [800]

Hoppe, Heidrun C., Benny Moldovanu, and Aner Sela (2009), "The theory of assortative matching based on costly signals." *Review of Economic Studies*, 76, 253–281. [802]

Kara, Tarik and Tayfun Sönmez (1996), "Nash implementation of matching rules." *Journal of Economic Theory*, 68, 425–439. [804]

Kemp, Adrienne W. (1998), "Absorption sampling and the absorption distribution." *Journal of Applied Probability*, 35, 489–494. [810, 820]

Rochet, Jean-Charles and Jean Tirole (2003), "Platform competition in two-sided markets." *Journal of the European Economic Association*, 1, 990–1029. [802]

Rochet, Jean-Charles and Jean Tirole (2006), "Two-sided markets: A progress report." *RAND Journal of Economics*, 37, 645–667. [802, 805, 816]

Roth, Alvin E. (1984a), "Misrepresentation and stability in the marriage problem." *Journal of Economic Theory*, 34, 383–387. [804]

Roth, Alvin E. (1984b), "The evolution of the labor market for medical interns and residents: A case study in game theory." *Journal of Political Economy*, 92, 991–1016. [808]

Rubinstein, Ariel and Asher Wolinsky (1987), "Middlemen." *Quarterly Journal of Economics*, 102, 581–593. [802]

Spiegler, Ran (2000), "Extracting interaction-created surplus." *Games and Economic Behavior*, 30, 142–162. [802]

Weyl, E. Glen (2010), "A price theory of multi-sided platforms." *American Economic Review*, 100, 1642–1672. [802]

Yavaş, Abdullah (1992), "Marketmakers versus matchmakers." *Journal of Financial Intermediation*, 2, 33–58. [802]