Journal of Mechanism and Institution Design
ISSN: 2399-844X (Print), 2399-8458 (Online)
DOI:10.22574/jmid.2024.12.005

RESPECT FOR IMPROVEMENTS AND COMPARATIVE STATICS IN MATCHING MARKETS

Scott Duke Kominers

Harvard University & a16z crypto, USA
kominers@fas.harvard.edu

ABSTRACT

One of the oldest results in the theory of two-sided matching is the *entry comparative static*, which shows that under the Gale–Shapley deferred acceptance algorithm, adding a new agent to one side of the market makes all the agents on the other side weakly better off. Here, we give a new proof of the entry comparative static, by way of a well-known property of deferred acceptance called *respect for improvements*. Our argument extends to yield comparative static results in more general settings, such as matching with slot-specific preferences.

Keywords: Matching, market entry, respect for improvements.

JEL Classification Numbers: C78, D47, D82.

I want to dedicate this article to Vincent Crawford – an extraordinary mentor, collaborator, and friend, who I learn from continuously (and also discretely). I appreciate the helpful comments of Péter Biró, Rachel Bayefsky, Vincent Crawford, Zoe DeStories, Piotr Dworczak, Federico Echenique, John William Hatfield, Emir Kamenica, Shengwu Li, Paul Milgrom, Michael Ostrovsky, Alvin Roth, Larry Samuelson, Tayfun Sönmez, Alexander Teytelboym, Bertan Turhan, Zaifu Yang, Bumin Yenmez, and especially Jiafeng Chen, Yunseo Choi, Jo Ellery, Ravi Jagadeesan, Shira Li, Jonathan Ma, Franklyn Wang, and other members of the Lab for Economic Design. I gratefully acknowledge the support of National Science Foundation grant SES-1459912 and the Ng Fund and the Mathematics in Economics Research Fund of the Harvard Center of Mathematical Sciences and Applications. Part of this work was conducted during the Simons Laufer Mathematical Sciences Institute Fall 2023 program on the Mathematics and Computer Science of Market and Mechanism Design, which was supported by the National Science Foundation under Grant No. DMS-1928930 and by the Alfred P. Sloan Foundation under grant G-2021-16778.

1. INTRODUCTION

IN their seminal 1962 paper, Gale & Shapley introduced the two-sided matching problem: We are given sets of men and women, each with preferences over members of the opposite gender; we seek a *stable matching*, i.e., an assignment of partners such that no one finds his or her partner unacceptable, and no two agents mutually prefer each other to their assigned partners.

Gale & Shapley (1962) showed that stable matchings can be found via the following *deferred acceptance* algorithm, under which men propose marriage in sequence, and women defer accepting prospective partners until the full sequence of proposals has been completed.

DEFERRED ACCEPTANCE ALGORITHM

Step 1**.** Each man proposes to his first-choice woman. Each woman *"holds"* her best acceptable proposal (if any) and rejects all other proposals.

Step $\ell \geq 2$**.** Each man who was rejected in the previous step proposes to his most-preferred woman to whom he has not yet proposed (if any). Each woman holds her best acceptable proposal among those made in this step and held from the previous step (if any) and rejects all others.

If at any time no men are available to make proposals—that is, if all men not currently held have proposed to all women they find acceptable—then the algorithm terminates. The *(man-proposing) deferred acceptance outcome* is the matching that matches each woman to the man whose proposal she is holding (if any) at the end of the last step before the algorithm terminates.

Variants of deferred acceptance are by now applied in a number of real-world contexts, including medical residency matching (Roth 1984; Roth & Peranson 1999), school choice (Balinski & Sönmez 1999; Abdulkadiroğlu & Sönmez 2003; Pathak 2011), and the assignment of cadets to branches of military service (Sönmez & Switzer 2013; Sönmez 2013; Greenberg et al. 2024).[1]

[1] For surveys of this and related work in *market design*, see, e.g., Roth (2008, 2013, 2015) and Kominers et al. (2017).

One of the most classic results in the theory of stable matching is an *entry comparative static* that characterizes how the deferred acceptance outcome changes when a new woman is added to the market (Kelso & Crawford 1982; Gale & Sotomayor 1985; Roth & Sotomayor 1990; Crawford 1991). The exact form of the comparative static varies across settings, but the core result is that *under deferred acceptance, adding a new agent to one side of the market makes all agents on the other side of the market weakly better off* (see, e.g.,Blum et al. 1997; Blum & Rothblum 2002; Hatfield & Milgrom 2005; Biró et al. 2008; Ostrovsky 2008; Hatfield & Kominers 2013; Chambers & Yenmez 2017; Yenmez 2018; Fleiner et al. 2018).[2]

Proofs of the entry comparative static and its generalizations tend to rely either on meticulous inductive arguments (Kelso & Crawford 1982; Crawford 1991; Blum et al. 1997; Blum & Rothblum 2002; Hatfield & Milgrom 2005; Biró et al. 2008; Hatfield & Kominers 2013; Fleiner et al. 2018) or on broad results characterizing the structure of the set of stable matchings (Gale & Sotomayor 1985; Ostrovsky 2008; Chambers & Yenmez 2017).[3]

In this paper, we give a new, more concise proof of the comparative static by way of another classic result—specifically, the fact that deferred acceptance *respects (unambiguous) improvements* in the sense that making one agent more highly ranked in other agents' preferences improves that agent's deferred acceptance outcome (Balinski & Sönmez 1999; Choi et al. 2024).[4] We show how to transform the entry of a new agent on one side of the market into a simultaneous preference rank improvement for the agents on the other side; the comparative static result then follows directly from respect for improvements.

As we describe, our approach extends naturally to yield novel comparative static results in more general settings, such as the "slot-specific preferences" framework that Kominers & Sönmez (2016) introduced to model matching with diversity constraints. In particular, we show that expanding the capacity

[2] As we discuss in Section 6, we typically also obtain a sort of dual to this result: agents on the same side of the market as the entering agent are made weakly worse off. To our knowledge, however, the methods we use here only enable us to derive the comparative static for agents on the side of the market opposite to that of the entering agent.

[3] Additionally, Choi (2022) used the logical compactness–based methods introduced by Gonczarowski et al. (2024) to establish a version of the entry comparative static for markets with infinitely many agents.

[4] See also the work of Wang & Kominers (2020), which uses the theory of matching with contracts to show a particularly strong link between respect for improvements and strategy-proofness.

of a firm that has slot-specific preferences weakly improves employees' match outcomes. Then, again in the slot-specific preference setting, we obtain what—to the best of our knowledge—are the first comparative statics for how the set of stable matchings changes when new contracts with an agent become possible.

2. AN ILLUSTRATIVE EXAMPLE

We consider a market with three men $\{\mathsf{Alex}, \mathsf{Bob}, \mathsf{Charlie}\}$ and two women $\{\mathsf{Xena}, \mathsf{Yvette}\}$, who rank each other as follows (with the option $\emptyset$ representing being unmatched):

$$\begin{aligned} \mathsf{Alex}&: \mathsf{Yvette} \succ \mathsf{Xena} \succ \emptyset & \mathsf{Xena}&: \mathsf{Alex} \succ \mathsf{Bob} \succ \mathsf{Charlie} \succ \emptyset \\ \mathsf{Bob}&: \mathsf{Yvette} \succ \mathsf{Xena} \succ \emptyset & \mathsf{Yvette}&: \mathsf{Alex} \succ \mathsf{Bob} \succ \emptyset. \\ \mathsf{Charlie}&: \mathsf{Xena} \succ \mathsf{Yvette} \succ \emptyset \end{aligned} \tag{1}$$

If we run deferred acceptance on our example, we see that first Alex and Bob propose to Yvette, while Charlie proposes to Xena; then Yvette and Xena hold Alex's and Charlie's proposals, respectively. Thus Bob is rejected; he proposes to Xena in the second step, causing Xena to hold Bob and reject Charlie. Finally, Charlie proposes to Yvette but is rejected, leaving us with an outcome μ in which Alex is matched to Yvette and Bob is matched to Xena. By construction, this matching μ is stable—everyone who is matched has an acceptable partner, and no man and woman mutually prefer each other to their assigned partners.

Now, we imagine that a new woman, Zelda, is added to the market we considered earlier, and that Zelda has the same preferences over men as Yvette. We suppose furthermore that Zelda is especially popular[5] so that preferences are now given by

$$\begin{aligned} \mathsf{Alex}&: \mathsf{Zelda} \succ \mathsf{Yvette} \succ \mathsf{Xena} \succ \emptyset & \mathsf{Xena}&: \mathsf{Alex} \succ \mathsf{Bob} \succ \mathsf{Charlie} \succ \emptyset \\ \mathsf{Bob}&: \mathsf{Zelda} \succ \mathsf{Yvette} \succ \mathsf{Xena} \succ \emptyset & \mathsf{Yvette}&: \mathsf{Alex} \succ \mathsf{Bob} \succ \emptyset \\ \mathsf{Charlie}&: \mathsf{Zelda} \succ \mathsf{Xena} \succ \mathsf{Yvette} \succ \emptyset & \mathsf{Zelda}&: \mathsf{Alex} \succ \mathsf{Bob} \succ \emptyset. \end{aligned} \tag{2}$$

Deferred acceptance now produces a matching $\tilde{\mu}$ under which Alex, Bob, and Charlie are matched to Zelda, Yvette, and Xena, respectively. Under $\tilde{\mu}$

[5] Perhaps she has an entire videogame franchise to her name.

every man is matched to a partner he weakly prefers to his partner under μ; the entry comparative static implies that this sort of improvement should arise in general.

Now, to illustrate our approach to the comparative static, we recall that deferred acceptance *respects improvements* in the sense that making one agent more highly ranked in other agents' preferences improves that agent's deferred acceptance outcome. We show in particular that the entry of a new woman such as Zelda can be modeled as a (weak) increase in the preference rank of every man in the market. Indeed, we could have written our original market (1) in the form

$$\begin{aligned} \mathsf{Alex} &: \mathsf{Zelda} \succ \mathsf{Yvette} \succ \mathsf{Xena} \succ \emptyset & \qquad \mathsf{Xena} &: \mathsf{Alex} \succ \mathsf{Bob} \succ \mathsf{Charlie} \succ \emptyset \\ \mathsf{Bob} &: \mathsf{Zelda} \succ \mathsf{Yvette} \succ \mathsf{Xena} \succ \emptyset & \qquad \mathsf{Yvette} &: \mathsf{Alex} \succ \mathsf{Bob} \succ \emptyset \\ \mathsf{Charlie} &: \mathsf{Zelda} \succ \mathsf{Xena} \succ \mathsf{Yvette} \succ \emptyset & \qquad \mathsf{Zelda} &: \emptyset \succ \mathsf{Alex} \succ \mathsf{Bob}, \end{aligned} \tag{3}$$

with Zelda treated as if she were already present, but found all men unacceptable. (Note that in (3) we have given Zelda a ranking of unacceptable men consistent with her preferences over men in (2), even though she will end up assigned to "$\emptyset$" rather than to any of them.) Now, the deferred acceptance outcome under (3) corresponds directly to that under (1): Alex is matched to Yvette and Bob is matched to Xena, just as under μ. Yet the "entry" of Zelda obtained by transforming (3) to (2) represents a ranking improvement for all the men, since each man is now more highly ranked relative to $\emptyset$. This implies by the respect for improvements result that all the men's deferred acceptance outcomes must weakly improve—precisely the claim of the entry comparative static.

We show in the sequel that the argument just described can be formalized to prove the full entry comparative static for the marriage model. Moreover, the argument extends directly to yield comparative static results in more general settings.

3. THE MARRIAGE MODEL

We start by introducing the marriage model of Gale & Shapley (1962): There are finite sets M and W of *men* and *women*; we denote by $I \equiv M \cup W$ the set of *agents*. We assume that each man $m \in M$ has a complete, transitive, and strict *preference ordering* $\succ_m$ over $W \cup \{\emptyset\}$, where $\emptyset$ denotes an *outside* *option* that represents the possibility of remaining unmatched. Similarly, each woman $w \in W$ has a complete, transitive, and strict *preference ordering* $\succ_w$ over $M \cup \{\emptyset\}$. For each agent $i \in I$, we denote the weak part of i's preferences by $\succcurlyeq_i$, so that if $j \succcurlyeq_i k$ then either $j \succ_i k$ or $j = k$. We use the convention that $\succ_{I'} \equiv (\succ_i)_{i \in I'}$; analogously, we write $\succcurlyeq_{I'} \equiv (\succcurlyeq_i)_{i \in I'}$.

3.1. Stable Matchings

A *(marriage) matching* is a map $\mu : I \to I \cup \{\emptyset\}$ from the set of agents to the set of agents plus the outside option, such that:

1. Under μ, each man is assigned to a woman or to the outside option—that is, $\mu(m) \in (W \cup \{\emptyset\})$ for each $m \in M$—and each woman is assigned to a man or to the outside option—$\mu(w) \in (M \cup \{\emptyset\})$ for each $w \in W$.

2. If a man $m \in M$ is assigned to a woman $w \in W$, then w is assigned to m—i.e., if $\mu(m) = w \in W$, then $\mu(w) = m$—and vice versa—if $\mu(w) = m \in M$, then $\mu(m) = w$.

We say that agent $i \in I$ is *matched to* $j \in I$ *under* μ if $\mu(i) = j$; we say that $i \in I$ is *unmatched under* μ if $\mu(i) = \emptyset$.

A matching μ is *individually rational* if no agent prefers the outside option to his or her assigned partner, i.e., if for all $i \in I$, we have that $\mu(i) \succcurlyeq_i \emptyset$. A matching μ is *unblocked* if there does not exist a man $m \in M$ and woman $w \in W$ who mutually prefer each other to their assigned match partners, i.e., such that $w \succ_m \mu(m)$ and $m \succ_w \mu(w)$. A matching μ is *stable* if it is both individually rational and unblocked.

The main result of Gale & Shapley (1962) shows that the deferred acceptance algorithm introduced in the Introduction always produces a stable outcome.

4. THE CLASSIC COMPARATIVE STATIC

Now, we show that the entry of a new woman $\tilde{w}$ to the market must make all the men weakly better off under (man-proposing) deferred acceptance. Abusing notation slightly, we extend our model to the set of agents $I \cup \{\tilde{w}\} = M \cup W \cup \{\tilde{w}\}$, writing $\tilde{\succ}$ and $\tilde{\mu}$, respectively, for preferences and matchings in the expanded market. To reflect the idea that $\tilde{w}$ is a new entrant to an existing market, we require that the preference profile $\tilde{\succ}$ be consistent with $\succ$ on I:

$$\text{for all } i, j, k \in I, \text{ we have } j \mathrel{\tilde{\succ}_i} k \iff j \succ_i k;$$

equivalently, for each man $m \in M$, the expanded preferences $\tilde{\succ}_m$ correspond to $\succ_m$ with $\tilde{w}$ added in somewhere, and for each woman $w \in W$, we have $\tilde{\succ}_w = \succ_w$.

Theorem 1 (Kelso & Crawford 1982; Gale & Sotomayor 1985; Roth & Sotomayor 1990)**.** *If μ^* is the outcome of man-proposing deferred acceptance in the market I and $\tilde{\mu}^*$ is the outcome of man-proposing deferred acceptance in the market $I \cup \{\tilde{w}\}$ that arises after the entry of woman $\tilde{w}$, then each man $m \in M$ (weakly) prefers his assignment under $\tilde{\mu}^*$ to his assignment under μ^*; that is,*

$$\tilde{\mu}^*(m) \mathrel{\tilde{\succsim}_m} \mu^*(m).[6] \tag{4}$$

One can prove Theorem 1 by explicitly tracking how the presence of $\tilde{w}$ affects each step of deferred acceptance (see, e.g., Crawford 1991; Chambers & Yenmez 2017). Another approach to Theorem 1 starts with the outcome of man-proposing deferred acceptance in the original market and then studies how the market re-equilibrates after $\tilde{w}$ enters (see, e.g., Kelso & Crawford 1982; Blum et al. 1997).[7]

Here, we give a novel proof of Theorem 1 that avoids the need to explicitly track either the algorithm or the path of adjustment. The key to our approach is a core property of deferred acceptance called *respect for (unambiguous) improvements*.

4.1. Respect for Improvements

For $w \in W$ and preference relations $\hat{\succ}_w$ and $\check{\succ}_w$, we say that $\hat{\succ}_w$ *unambiguously improves upon* $\check{\succ}_w$ *for* $m \in M$ if m is ranked (weakly) higher under $\hat{\succ}_w$ than under $\check{\succ}_w$; the relative rankings of all other men are left unchanged; and no man becomes unacceptable where he was previously acceptable. Formally:

6 Theorem 1 is also a special case of the comparative static results of Crawford (1991); Hatfield & Milgrom (2005); Ostrovsky (2008); Hatfield & Kominers (2013); Chambers & Yenmez (2017); Yenmez (2018); Fleiner et al. (2018).

7 Dworczak (2021) tracks similar re-equilibration dynamics in his *Deferred Acceptance with Compensation Chains* algorithms, which generalize deferred acceptance by allowing agents on both sides of the market to make proposals.

Definition 1. We say that $\hat{\succ}_w$ is an *unambiguous improvement over* $\check{\succ}_w$ *for* $m \in M$ if

- for all $m' \in ((M \setminus \{m\}) \cup \{\emptyset\})$, if $m \check{\succ}_w m'$ then $m \hat{\succ}_w m'$;
- for all $m', m'' \in (M \setminus \{m\})$, we have $m' \check{\succ}_w m''$ if and only if $m' \hat{\succ}_w m''$; and
- for all $m' \in (M \setminus \{m\})$, if $m' \check{\succ}_w \emptyset$ then $m' \hat{\succ}_w \emptyset$.

We say that a profile of the women's preferences $\hat{\succ}_W$ is an *unambiguous improvement over* $\check{\succ}_W$ *for* m if each $\hat{\succ}_w$ is an unambiguous improvement over $\check{\succ}_w$ for m.

Balinski & Sönmez (1999) showed that deferred acceptance *respects (unambiguous) improvements* in the sense that if $\hat{\succ}_W$ is an unambiguous improvement over $\check{\succ}_W$ for m, then m (weakly) prefers his man-proposing deferred acceptance outcome when the women's preferences are $\hat{\succ}_W$ to that when the women's preferences are $\check{\succ}_W$.[8]

Lemma 1. *If* $\hat{\succ}_W$ *is an unambiguous improvement over* $\check{\succ}_W$ *for m, and* $\hat{\mu}^*$ *and* $\check{\mu}^*$ *denote the man-proposing deferred acceptance outcomes under* $(\succ_M, \hat{\succ}_W)$ *and* $(\succ_M, \check{\succ}_W)$*, respectively, then m (weakly) prefers his or her assignment under* $\hat{\mu}^*$ *to that under* $\check{\mu}^*$*; that is,*

$$\hat{\mu}^*(m) \succcurlyeq_m \check{\mu}^*(m).$$

Respect for unambiguous improvements is an extremely natural property to desire in practice, as it means that agents have no incentive to try to lower their standing in others' preference relations.[9,10] But the respect for improvements condition is also subtle: If a man's rank just improves relative to the outside option $\emptyset$, then the second and third conditions of Definition 1 hold automatically. Thus a (strict) unambiguous improvement for one man can also

[8] As discussed by Choi et al. (2024), the form of the respect for unambiguous improvements condition stated in Definition 1 is slightly weaker than the version used by Balinski & Sönmez (1999); as a result, the statement of Lemma 1 is slightly stronger than the Balinski & Sönmez (1999) version of the respect for improvements result. But the essence of the result is the same—and formally, our Lemma 1 follows from Lemma 2, which in turn comes from Choi et al. (2024).

[9] Imagine the chaos if intentional self-sabotage were a standard feature of the marriage market!

[10] Moreover, respect for unambiguous improvements in some sense uniquely characterizes deferred acceptance (see Theorems 5 and 6 of Balinski & Sönmez 1999).

be a (strict) unambiguous improvement for some other man, so long as both men's positions only improve relative to the outside option $\emptyset$; this property allows us to derive an elegant and efficient proof of Theorem 1.[11]

4.2. Proof of Theorem 1

As in the statement of Theorem 1, we suppose that μ^* is the outcome of man-proposing deferred acceptance in the market I and $\tilde{\mu}^*$ is the outcome of man-proposing deferred acceptance in the market $I \cup \{\tilde{w}\}$.

We write the preferences of $\tilde{w}$ in the form

$$\tilde{\succ}_{\tilde{w}} : m_1 \,\tilde{\succ}_{\tilde{w}} \cdots \tilde{\succ}_{\tilde{w}}\, m_\ell \,\tilde{\succ}_{\tilde{w}}\, \emptyset \,\tilde{\succ}_{\tilde{w}}\, m_{\ell+1} \cdots \tag{5}$$

and let $\bar{\succ}_{\tilde{w}}$ be the "null" preference relation for $\tilde{w}$ that is consistent with $\tilde{\succ}_{\tilde{w}}$ but treats all men as unacceptable:

$$\bar{\succ}_{\tilde{w}} : \emptyset \,\bar{\succ}_{\tilde{w}}\, m_1 \,\bar{\succ}_{\tilde{w}} \cdots \bar{\succ}_{\tilde{w}}\, m_\ell \,\bar{\succ}_{\tilde{w}}\, m_{\ell+1} \cdots . \tag{6}$$

We let $\bar{\mu}^*$ be the outcome of man-proposing deferred acceptance in the market with agents $I \cup \{\tilde{w}\}$ and preference profile $(\tilde{\succ}_M, \tilde{\succ}_W, \bar{\succ}_{\tilde{w}})$. As $\tilde{\succ}_W = \succ_W$ and each $\tilde{\succ}_m$ agrees with $\succ_m$ except for the inclusion of $\tilde{w}$, we see directly that

$$\bar{\mu}^*(i) = \mu^*(i) \tag{7}$$

for each $i \in I$. Indeed, as $\tilde{w}$ treats all men as unacceptable under $\bar{\succ}_{\tilde{w}}$, deferred acceptance's outcome is unchanged if all men drop $\tilde{w}$ from their preference relations—and with that preference adjustment, the algorithm corresponds exactly to deferred acceptance in the market I, and so must yield an outcome that corresponds to μ^*.

Now, we note by inspecting (5) and (6) that $\tilde{\succ}_{\tilde{w}}$ is an unambiguous improvement over $\bar{\succ}_{\tilde{w}}$ for each man $m \in M$. Hence, by Lemma 1, we see that each man $m \in M$ (weakly) prefers his assignment under $\tilde{\mu}^*$ to that under $\bar{\mu}^*$; that is,

$$\tilde{\mu}^*(m) \,\tilde{\succeq}_m\, \bar{\mu}^*(m) \tag{8}$$

for each man $m \in M$. Combining (7) and (8) shows the desired result, (4).

[11] As this intuition suggests, the argument does not require the full strength of the respect for improvements result—we use only respect for improvements relative to the outside option $\emptyset$.

4.3. Discussion

Our proof of Theorem 1 highlights that the key feature underlying Theorem 1 is that $\tilde{w}$'s entry simultaneously improves all the men's overall standing in the market by increasing competition among the women. This idea is closely related to a much earlier argument due to Gale & Sotomayor (1985). Like us, Gale & Sotomayor (1985) treated $\tilde{w}$'s entry as equivalent to transforming $\tilde{w}$'s preferences from a null relation to $\tilde{\succsim}_{\tilde{w}}$. Unlike in our argument, however, Gale & Sotomayor (1985) derived the comparative static through appeal to a result showing that the set of stable outcomes has a lattice structure (see Roth & Sotomayor 1990, pp. 43–44). Approaching the argument in terms of respect for improvements instead is useful because it immediately suggests substantial generalizations, some of which are outside the scope of the lattice structure result—as we describe next.

5. A MORE GENERAL MODEL: MATCHING WITH SLOT-SPECIFIC PREFERENCES

We now extend Theorem 1 to a setting with *many-to-one matching*—that is, one in which employees on one side of the market may take multiple partners on the other side (Gale & Shapley 1962; Roth 1985; Roth & Sotomayor 1990). We furthermore generalize by allowing the matching process to determine not just who matches with whom but also *contracts* that specify terms of exchange like wages or hours worked (Crawford & Knoer 1981; Kelso & Crawford 1982; Hatfield & Milgrom 2005). Specifically, we work with the *slot-specific preference* structure introduced by Kominers & Sönmez (2016).

5.1. Intuition

The slot-specific preferences framework, which we describe formally in the next section, is a model of employee–firm matching in which each firm has a set of positions—*slots*—that can be assigned to different employees. Slots have their own (potentially independent) rankings over contracts. Within each firm, a linear order called the *order of precedence* determines the order in which slots are filled.

For example, we might imagine a firm with two slots, s^1 and s^2, one of which can be used to hire any worker at either a high ($£_H$) or low ($£_L$) salary, while the second slot can only be used to hire workers at a lower salary:

$$s^1 : (i, £_L) \succ (i, £_H) \succ (j, £_L) \succ (j, £_H) \succ \cdots$$
$$s^2 : (i, £_L) \succ (j, £_L) \succ \cdots .$$

One special case of slot-specific preferences is when we make all the slots within a given firm identical, making that firm's preferences consistent with a single linear order over contracts and a maximum number of positions that can be filled (the *responsive preference* model of Roth 1985). More broadly, slot-specific preferences embed multiple models of affirmative action, such as using some slots to reserve positions for members of disadvantaged groups (Abdulkadiroğlu 2005; Kojima 2012; Hafalir et al. 2013; Kominers & Sönmez 2014; Dur et al. 2018, 2020).

5.2. Formal Model

We suppose that there is a set of *employees* E and a set of *firms* F, and a (finite) set of *contracts* X. Each contract $x \in X$ is between an employee $\mathsf{e}(x) \in E$ and firm $\mathsf{f}(x) \in F$, and may also specify additional "terms" of exchange drawn from a set T. Thus X may be considered a subset of $E \times F \times T$. We extend the notations $\mathsf{e}(\cdot)$ and $\mathsf{f}(\cdot)$ to sets of contracts $Y \subseteq X$ by setting $\mathsf{e}(Y) \equiv \cup_{y \in Y} \{\mathsf{e}(y)\}$ and $\mathsf{f}(Y) \equiv \cup_{y \in Y} \{\mathsf{f}(y)\}$. For $Y \subseteq X$, we denote $Y_e \equiv \{y \in Y : \mathsf{e}(y) = e\}$ and $Y_{E'} \equiv \cup_{e \in E'} Y_e$; analogously, we denote $Y_f \equiv \{y \in Y : \mathsf{f}(y) = f\}$ and $Y_{F'} \equiv \cup_{f \in F'} Y_f$.

Each employee $e \in E$ has a complete, transitive, and strict preference order $\succ_e$ (with weak order $\succsim_e$) over contracts in $X_e \cup \{\emptyset\}$, where, as before, $\emptyset$ is an outside option that represents remaining unmatched; we use the convention that $\emptyset \succ_e x$ for all $x \in X \setminus X_e$. We say that the contracts $x \in X$ for which $\emptyset \succ_e x$ are *unacceptable to e*. We denote the profile of all employees' preferences by $\succ_E$.

Each firm $f \in F$ has a set S_f of *slots*; each slot can be assigned up to one contract in X_f. Slots $s \in S_f$ have (linear) preference orders $\succ_s$ (with weak orders $\succsim_s$) over contracts in X_f. As with employees, we assume that each slot $s \in S_f$ ranks an outside option $\emptyset$ that represents remaining unassigned, and as with employees, we use the convention that $\emptyset \succ_s x$ if $x \in X \setminus X_f$. We set $S \equiv \cup_{f \in F} S_f$ and denote the profile of all slots' preferences by $\succ_S$.

Employees have *unit demand*, that is, they choose at most one contract from a set of contract offers. We assume also that employees always choose the best available contract, so that the choice $C^e(Y)$ of an employee $e \in E$ from contract set $Y \subseteq X$ is the $\succ_e$-maximal element of Y_e (or the outside option if $\emptyset \succ_e y$ for all $y \in Y_e$).[12]

Meanwhile, firms $f \in F$ may be assigned as many as $q_f \equiv |S_f|$ contracts—one for each slot in S_f—but may hold no more than one contract with a given employee. We assume that for each $f \in F$, the slots in S_f are ordered according to a (linear) *order of precedence* $\rhd^f$. We denote $S_f \equiv \left\{s_f^1, \ldots, s_f^{q_f}\right\}$ with the understanding that $s_f^\ell \rhd^f s_f^{\ell+1}$ unless otherwise noted. The interpretation of $\rhd^f$ is that if $s \rhd^f s'$ then—whenever possible—firm f fills slot s before filling s'.

Formally, the choice $C^f(Y)$ of a firm $f \in F$ from contract set $Y \subseteq X$ is defined as follows:

- First, slot s_f^1 is assigned the contract x^1 that is $\succ_{s_f^1}$-maximal among contracts in Y.
- Then, slot s_f^2 is assigned the contract x^2 that is $\succ_{s_f^2}$-maximal among contracts in the set $Y \setminus Y_{\mathsf{e}(x^1)}$ of contracts in Y with employees other than $\mathsf{e}(x^1)$.
- This process continues in sequence, with each slot s_f^ℓ being assigned the contract x^ℓ that is $\succ_{s_f^\ell}$-maximal among contracts in the set

$$Y \setminus Y_{\mathsf{e}(\{x^1,\ldots,x^{\ell-1}\})}.$$

If no contract $x \in Y$ is assigned to slot $s_f^\ell \in S_f$ in the computation of $C^f(Y)$, then s_f^ℓ is assigned the null contract $\emptyset$.

5.3. Stable Outcomes

An *outcome* is a set of contracts $Y \subseteq X$ that is "feasible" in the sense that

- Y contains at most one contract for each employee, i.e., $|Y_e| \leq 1$ for each $e \in E$, and
- Y contains at most q_f contracts for each firm f, i.e., $|Y_e| \leq q_f$ for each $f \in F$.

[12] To simplify our exposition and notation, we treat individual contracts as interchangeable with singleton contract sets.

We say that an outcome Y is *stable* if it is

1. *individually rational*—$C^e(Y) = Y_e$ for all $e \in E$ and $C^f(Y) = Y_f$ for all $f \in F$—and
2. *unblocked*—there does not exist a firm $f \in F$ and *blocking set* $Z \neq C^f(Y)$ such that $Z = C^f(Y \cup Z)$ and $Z_e = C^e(Y \cup Z)$ for all $e \in \mathsf{e}(Z)$.

Note that if all employee–firm pairs may contract and the set T of contractual terms is trivial, then $X = E \times F$. If moreover all workers and firms have unit demand (i.e., if we always have $|C^e(Y)| \leq 1$ and $|C^f(Y)| \leq 1$), then we recover the marriage model from Section 3, where the man–woman pairs matched under some matching μ correspond to the pairs contained in an outcome Y.

5.4. Generalized Deferred Acceptance

Kominers & Sönmez (2016) showed that stable outcomes exist under slot-specific preferences, and can be found via the following *(employee-proposing) cumulative offer process* (Kelso & Crawford 1982; Hatfield & Milgrom 2005), which generalizes deferred acceptance:[13]

CUMULATIVE OFFER PROCESS

Step 1. Some employee $e^1 \in E$ proposes their most-preferred acceptable contract, $x^1 \in X_{e^1}$. Firm $\mathsf{f}(x^1)$ *holds* x^1 if $x^1 \in C^{\mathsf{f}(x^1)}(\{x^1\})$, and *rejects* x^1 otherwise. Set $A^2_{\mathsf{f}(x^1)} = \{x^1\}$, and set $A^2_{f'} = \varnothing$ for each $f' \neq \mathsf{f}(x^1)$; these are the sets of contracts *available* to firms at the beginning of Step 2.

Step $\ell \geq 2$. Some employee $e^\ell \in E$ for whom no contract is currently held by any firm proposes their most-preferred acceptable contract that has not yet been rejected, $x^\ell \in X_{e^\ell}$.

[13] The cumulative offer process generalizes deferred acceptance by (in principle) allowing firms to hold contracts they had rejected in earlier steps—although under slot-specific preferences, firms never actually use this extra degree of freedom (see, e.g., Hatfield et al. 2021). For consistency with Kominers & Sönmez (2016), we state the cumulative offer process with a single agent proposing in each step; Hirata & Kasuya (2014) have proven that in our setting this formulation is equivalent to one in which employees propose simultaneously, directly generalizing the version of deferred acceptance we presented in the Introduction.

> Firm $\mathsf{f}(x^\ell)$ holds the contracts in $C^{\mathsf{f}(x^\ell)}(A^\ell_{\mathsf{f}(x^\ell)} \cup \{x^\ell\})$ and rejects all other contracts in $A^\ell_{\mathsf{f}(x^\ell)} \cup \{x^\ell\}$; firms $f' \neq \mathsf{f}(x^\ell)$ continue to hold all contracts they held at the end of Step $\ell - 1$. Set $A^{\ell+1}_{\mathsf{f}(x^\ell)} = A^\ell_{\mathsf{f}(x^\ell)} \cup \{x^\ell\}$, and set $A^{\ell+1}_{f'} = A^\ell_{f'}$ for each $f' \neq \mathsf{f}(x^\ell)$.

If at any time no employee is able to propose a new contract—that is, if all employees for whom no contracts are on hold have proposed all the contracts they find acceptable—then the algorithm terminates. The *outcome of the (employee-proposing) cumulative offer process* is the set of contracts held by firms at the end of the last step before the algorithm terminates.[14]

Although stable outcomes exist under slot-specific preferences, the set of stable outcomes under slot-specific preferences does not have lattice structure (see Kominers & Sönmez 2016; Hatfield & Kominers 2020). Consequently the approach Gale & Sotomayor (1985) used to prove Theorem 1 does not carry over to the slot-specific preference setting. Nevertheless, as we show next, our approach based on respect for improvements extends directly.

5.5. Respect for Improvements Under Slot-Specific Preferences

Following Kominers & Sönmez (2016) and Choi et al. (2024), $\hat{\succ}_S$ *is an unambiguous improvement over preference profile* $\check{\succ}_S$ *for* $e \in E$ if $\hat{\succ}_S$ is obtained from $\check{\succ}_S$ by raising the positions of some of e's contracts (at some slots) while leaving the relative preference orders of other employees' contracts unchanged (and not making any contracts that were previously acceptable at some slot unacceptable[15]). Formally:

Definition 2. We say that preference profile $\hat{\succ}_S$ is an *unambiguous improvement over* $\check{\succ}_S$ *for* $e \in E$ if for all slots $s \in S$:

[14] Here we refer to "the" outcome of the cumulative offer process because—at least under slot-specific preferences—the outcome is independent of the order proposals (see Kominers & Sönmez 2016, as well as Hirata & Kasuya 2014).

[15] This last condition—which corresponds to the third condition in Definition 2—was not included by Kominers & Sönmez (2016); however, as Choi et al. (2024) showed, correcting Kominers & Sönmez (2016), it is actually needed for the respect for improvements result to work as stated.

1. for all $x \in X_e$ and $y \in (X_{E \setminus \{e\}} \cup \{\emptyset\})$, if $x \mathrel{\check{\succ}_s} y$, then $x \mathrel{\hat{\succ}_s} y$;
2. for all $y, z \in X_{E \setminus \{e\}}$, $y \mathrel{\hat{\succ}_s} z$ if and only if $y \mathrel{\check{\succ}_s} z$; and
3. for all $y \in X_{E \setminus \{e\}}$, if $y \mathrel{\check{\succ}_s} \emptyset$, then $y \mathrel{\hat{\succ}_s} \emptyset$.

The cumulative offer process *respects unambiguous improvements* in the sense that if $\hat{\succ}_S$ is an unambiguous improvement over $\check{\succ}_S$ for e, then e (weakly) prefers their cumulative offer process outcome under $(\succ_E, \hat{\succ}_S)$ to that under $(\succ_E, \check{\succ}_S)$.

Lemma 2 (Kominers & Sönmez 2016; Choi et al. 2024)**.** *If $\hat{\succ}_S$ is an unambiguous improvement over $\check{\succ}_S$ for e, and $\hat{Y}^*$ and $\check{Y}^*$ denote the employee-proposing cumulative offer process outcomes under $(\succ_E, \hat{\succ}_S)$ and $(\succ_E, \check{\succ}_S)$, respectively, then e (weakly) prefers their assignment under $\hat{Y}^*$ to that under $\check{Y}^*$; that is,*

$$\hat{Y}^*_e \succsim_e \check{Y}^*_e.$$

5.6. Comparative Statics

Lemma 2 immediately implies several comparative statics for markets with slot-specific preferences—generalizing Theorem 1—through a version of the argument we presented in Section 4.2.

5.6.1. Expanding Capacity

First, we consider what happens to the cumulative offer process outcome when we add a slot $\tilde{s}$ at firm f. Abusing notation again, we extend our slot-specific preference model to the set of slots $S \cup \{\tilde{s}\}$, writing $\tilde{\succ}_S$ and $\tilde{Y}$, respectively, for slot preferences and outcomes in the expanded market ($\tilde{s}$ can appear anywhere in the precedence order). To reflect the idea that $\tilde{s}$ is a new addition to an existing market, we require that preferences of slots in S are unchanged: $\tilde{\succ}_S = \succ_S$. (Note that here we may also leave the employee preferences $\succ_E$ unchanged, as we have not in any way changed the set of firms F or the set of contracts X.)

By construction, adding the new slot $\tilde{s}$ impacts the market exactly the way that adding a new woman impacted the market in Section 4.2: it is as if we raised the position of all agents' contracts at a slot that previously treated all contracts as unacceptable, holding all other slots' preferences fixed. Hence, adding $\tilde{s}$ results in an unambiguous improvement for all agents; this implies the following comparative static.

Theorem 2. *If Y^* is the outcome of the employee-proposing cumulative offer process in the market with the set of slots S and $\tilde{Y}^*$ is the outcome of the employee-proposing cumulative offer process in the market with the set of slots $S \cup \{\tilde{s}\}$, then each employee $e \in E$ (weakly) prefers their assignment under $\tilde{Y}^*$ to their assignment under Y^*; that is,*

$$\tilde{Y}^*_e \succeq_e Y^*_e.$$

Theorem 2 generalizes Theorem 1, as in the slot-specific preferences framework we can model entry of a new woman into a marriage market as adding a single slot at a firm that previously had none. Theorem 2 also has practical implications: Perhaps most naturally, the result means that under (employee-proposing) deferred acceptance, expanding the number of available positions at a firm always works in employees' favor. Moreover, by applying Theorem 2 iteratively (adding one slot at a time), we obtain the conclusion of Theorem 2 for the entry of a wholly new employer to the market.

5.6.2. Adding Contracts

A further generalization of Theorem 1 shows that adding new contracts at the bottom of some slots' preference orders (that is, right before the null contracts $\emptyset$) again results in improved outcomes for all employees.

Suppose that we introduce new contracts $\tilde{X}$, yielding the contract set $X \cup \tilde{X}$. We write $\tilde{\succ}$ and $\tilde{Y}$ for preferences and outcomes in the expanded market, assuming that $y \mathrel{\tilde{\succ}_r} y'$ if and only if $y \succ_r y'$ for all $r \in E \cup S$ and $y, y' \in X$—reflecting the idea that adding $\tilde{x}$ does not affect agents' or slots' preferences over pre-existing contract options. If $y \mathrel{\tilde{\succ}_s} \tilde{x}$ for all $s \in S$, $y \in X$, and $\tilde{x} \in \tilde{X}$, then, just as in the argument we used in Section 4.2, we may interpret $\tilde{\succ}_S$ as an unambiguous improvement over $\succ_S$ under the contract set $X \cup \tilde{X}$ by imagining that $\succ_S$ ranks all the contracts in $\tilde{X}$ as unacceptable—with relative ranking consistent with $\tilde{\succ}_S$—so that $\tilde{\succ}_S$ can be obtained from $\succ_S$ by raising the position of contracts in $\tilde{X}$ relative to the outside option. Thus, we have the following comparative static by Lemma 2.

Theorem 3. *If Y^* is the outcome of the employee-proposing cumulative offer process in the market with the set of contracts X and $\tilde{Y}^*$ is the outcome of* *the employee-proposing cumulative offer process in the market with the set of contracts $X \cup \tilde{X}$ (with $y \,\tilde{\succ}_s\, \tilde{x}$ for all $s \in S$, $y \in X$, and $\tilde{x} \in \tilde{X}$), then each employee $e \in E$ (weakly) prefers their assignment under $\tilde{Y}^*$ to their assignment under Y^*; that is,*

$$\tilde{Y}^*_e \,\tilde{\succsim}_e\, Y^*_e.$$

Lastly, we note that adding new contracts for a single employee $e \in E$ *anywhere* in slots' preferences also results in an unambiguous improvement—and hence ensures that e will be better off under the cumulative offer process.[16]

Indeed, suppose that we introduce a new contract $\tilde{x}$, yielding the contract set $X \cup \{\tilde{x}\}$. We write $\tilde{\succ}$ and $\tilde{Y}$ for preferences and outcomes in the expanded market, assuming that $y \,\tilde{\succ}_r\, y'$ if and only if $y \succ_r y'$ for all $r \in E \cup S$ and $y, y' \in X$—again reflecting the idea that adding $\tilde{x}$ does not affect agents' or slots' preferences over pre-existing contract options. By construction, adding the new contract $\tilde{x}$ results in an unambiguous improvement for $\mathsf{e}(\tilde{x})$: employee $\mathsf{e}(\tilde{x})$ has weakly–higher ranked contracts at every slot, while relative rankings of all other contracts are unchanged; this implies the following comparative static by Lemma 2.

Theorem 4. *If Y^* is the outcome of the employee-proposing cumulative offer process in the market with the set of contracts X and $\tilde{Y}^*$ is the outcome of the employee-proposing cumulative offer process in the market with the set of contracts $X \cup \{\tilde{x}\}$, then $\mathsf{e}(\tilde{x})$ (weakly) prefers their assignment under $\tilde{Y}^*$ to their assignment under Y^*; that is,*

$$\tilde{Y}^*_{\mathsf{e}(\tilde{x})} \,\tilde{\succsim}_{\mathsf{e}(\tilde{x})}\, Y^*_{\mathsf{e}(\tilde{x})}.$$

5.7. Discussion

Our use of the respect for improvements result (Lemma 2) here is similar to an early application of Kominers & Sönmez (2014) for the slot-specific preferences framework. Indeed, Kominers & Sönmez (2014) used respect for improvements under slot-specific preferences to show that guaranteeing slots at a school for minorities (Hafalir et al. 2013) improves welfare relative to simply capping the number of majority students allowed to attend that school (Kojima 2012). The crux of the Kominers & Sönmez (2014) argument consists

[16] We cannot in general add contracts at arbitrary positions for multiple agents, however, as the resulting change in preferences might not be an unambiguous improvement.

of observing that converting a quota slot into a reserve slot is an unambiguous improvement for all majority students, as it corresponds to raising majority students' positions relative to the null option at each quota slot (while still ranking majority students below minority students at those slots); this is analogous to the way we apply Lemma 2.

Meanwhile, Chambers & Yenmez (2017) showed a result analogous to Theorem 2 for ways of "expanding" choice rules that satisfy a regularity condition called *path-independence* (Aizerman & Malishevski 1981). Yenmez (2018) extended the Chambers & Yenmez (2017) result still further to cover choice rules that can be modified—or in the language of Hatfield & Kominers (2020), *completed*—in ways that make them path-independent (see also Echenique & Yenmez 2015; Kamada & Kojima 2015). Slot-specific preferences are not path-independent, in general, but *do* have path-independent completions (see Hatfield & Kominers 2020). However, the additions of slots and contracts considered in Theorems 2 to 4 are *not* expansions in the sense of Chambers & Yenmez (2017) and Yenmez (2018); hence, our results here extend Chambers & Yenmez's (2017) and Yenmez's (2018) comparative statics to new types of transformations. To our knowledge, we are the first to prove comparative statics for the addition of new contracts to the market.[17]

6. CONCLUSION

Using the *respect for improvements* property of deferred acceptance, we have developed a new method of proving entry comparative statics in matching markets. Our method generalizes readily to any matching setting for which a respect for improvements result is known: We illustrated one such generalization in Section 5, where we used respect for improvements to show comparative statics for matching settings with slot-specific preferences.[18,19]

[17] That said, a few comparative statics have been shown previously for structured changes to the full set of contracts, such as adjustments to the level at which transfers between agents are taxed (see, e.g., Dupuy et al. 2020).

[18] Likewise, we can extend the entry comparative static to some matching settings under weakened substitutability conditions using the respect for improvements results of Afacan (2017, 2022).

[19] It also seems possible that variations of the standard respect for improvements concept, such as Choi & Li's (2024) *respect for improvements with manipulations*, would give rise to corresponding versions of the entry comparative static.

We note, however, that others have found more refined comparative static results through appeal to structural results for the set of stable matchings. Most common is a sort of dual to Theorem 1, showing that entry of a new woman to the market makes all the other women (weakly) worse off (see, e.g., Gale & Sotomayor 1985; Roth & Sotomayor 1990; Crawford 1991). Others have given more precise characterizations of which agents are helped and/or harmed by other agents' entry (see, e.g., Romm 2014), or shown comparative statics for the full set of stable matchings, rather than just for the deferred acceptance outcome (see, e.g., Blum et al. 1997; Chambers & Yenmez 2017). It is unclear whether our respect for improvements–based approach can also be used to prove the refinements just described. At a minimum, such an argument would seem to require a sharper version of the respect for improvements theorem. If we may pun a bit, let us leave such further "improvements" for future work.

References

Abdulkadiroğlu, A. (2005). College admissions with affirmative action. *International Journal of Game Theory*, *33*, 535–549.

Abdulkadiroğlu, A., & Sönmez, T. (2003). School choice: A mechanism design approach. *American Economic Review*, *93*(3), 729–747.

Afacan, M. O. (2017). Some further properties of the cumulative offer process. *Games and Economic Behavior*, *104*, 656–665.

Afacan, M. O. (2022). Corrigendum to "Some further properties of the cumulative offer process" [Games Econ. Beh. 104 (2017) 656–665]. *Games and Economic Behavior*, *133*, 352–354.

Aizerman, M. A., & Malishevski, A. V. (1981). General theory of best variants choice: Some aspects. *IEEE Transactions on Automatic Control*, *26*(5), 1030–1040.

Balinski, M., & Sönmez, T. (1999). A tale of two mechanisms: Student placement. *Journal of Economic Theory*, *84*(1), 73–94.

Biró, P., Cechlárová, K., & Fleiner, T. (2008). The dynamics of stable matchings and half-matchings for the stable marriage and roommates problems. *International Journal of Game Theory*, *36*(3-4), 333–352.

Blum, Y., Roth, A. E., & Rothblum, U. G. (1997). Vacancy chains and equilibration in senior-level labor markets. *Journal of Economic Theory*, *76*(2), 362–411.

Blum, Y., & Rothblum, U. G. (2002). "Timing is everything" and marital bliss. *Journal of Economic Theory*, *103*(2), 429–443.

Chambers, C. P., & Yenmez, M. B. (2017). Choice and matching. *American Economic Journal: Microeconomics*, *9*(3), 126–147.

Choi, Y. (2022). *On two-sided matching in infinite markets.* (Harvard University Working Paper)

Choi, Y., Kominers, S. D., & Sönmez, T. (2024). *Matching with slot-specific priorities: Theory – corrigendum.* (In Preparation)

Choi, Y., & Li, S. (2024). *Respect for improvements with manipulations.* (Harvard University Working Paper)

Crawford, V. P. (1991). Comparative statics in matching markets. *Journal of Economic Theory*, *54*(2), 389–400.

Crawford, V. P., & Knoer, E. M. (1981). Job matching with heterogeneous firms and workers. *Econometrica*, *49*(2), 437–450.

Dupuy, A., Galichon, A., Jaffe, S., & Kominers, S. D. (2020). Taxation in matching markets. *International Economic Review*, *61*(4), 1591–1634.

Dur, U., Kominers, S. D., Pathak, P. A., & Sönmez, T. (2018). Reserve design: Unintended consequences and the demise of Boston's walk zones. *Journal of Political Economy*, *126*(6), 2457–2479.

Dur, U., Pathak, P. A., & Sönmez, T. (2020). Explicit vs. statistical targeting in affirmative action: Theory and evidence from Chicago's exam schools. *Journal of Economic Theory*, *187*, 104996.

Dworczak, P. (2021). Deferred acceptance with compensation chains. *Operations Research*, *69*(2), 456–468.

Echenique, F., & Yenmez, M. B. (2015). How to control controlled school choice. *American Economic Review*, *105*(8), 2679–2694.

Fleiner, T., Jankó, Z., Tamura, A., & Teytelboym, A. (2018). *Trading networks with bilateral contracts.* (Oxford University Working Paper)

Gale, D., & Shapley, L. S. (1962). College admissions and the stability of marriage. *American Mathematical Monthly*, *69*(1), 9–15.

Gale, D., & Sotomayor, M. A. O. (1985). Some remarks on the stable matching problem. *Discrete Applied Mathematics*, *11*(3), 223–232.

Gonczarowski, Y. A., Kominers, S. D., & Shorrer, R. I. (2024). To infinity and beyond: A general framework for scaling economic theories. *Forthcoming in Theoretical Economics*.

Greenberg, K., Pathak, P. A., & Sönmez, T. (2024). Redesigning the US Army's branching process: A case study in minimalist market design. *American Economic Review*, *114*(4), 1070–1106.

Hafalir, I. E., Yenmez, M. B., & Yildirim, M. A. (2013). Effective affirmative action in school choice. *Theoretical Economics*, *8*(2), 325–363.

Hatfield, J. W., & Kominers, S. D. (2013). Vacancies in supply chain networks. *Economics letters*, *119*(3), 354–357.

Hatfield, J. W., & Kominers, S. D. (2020). *Hidden substitutes.* (Harvard University

Working Paper)

Hatfield, J. W., Kominers, S. D., & Westkamp, A. (2021). Stability, strategy-proofness, and cumulative offer mechanisms. *Review of Economic Studies*, *88*(3), 1457–1502.

Hatfield, J. W., & Milgrom, P. (2005). Matching with contracts. *American Economic Review*, *95*(4), 913–935.

Hirata, D., & Kasuya, Y. (2014). Cumulative offer process is order-independent. *Economics Letters*, *124*(1), 37–40.

Kamada, Y., & Kojima, F. (2015). Efficient matching under distributional constraints: Theory and applications. *American Economic Review*, *105*, 67–99.

Kelso, A. S., Jr., & Crawford, V. P. (1982). Job matching, coalition formation, and gross substitutes. *Econometrica*, *50*(6), 1483–1504.

Kojima, F. (2012). School choice: Impossibilities for affirmative action. *Games and Economic Behavior*, *75*, 685–693.

Kominers, S. D., & Sönmez, T. (2014). *Designing for diversity in matching.* (Harvard University Working Paper)

Kominers, S. D., & Sönmez, T. (2016). Matching with slot-specific priorities: Theory. *Theoretical Economics*, *11*(2), 683–710.

Kominers, S. D., Teytelboym, A., & Crawford, V. P. (2017). An invitation to market design. *Oxford Review of Economic Policy*, *33*(4), 541–571.

Ostrovsky, M. (2008). Stability in supply chain networks. *American Economic Review*, *98*(3), 897–923.

Pathak, P. A. (2011). The mechanism design approach to student assignment. *Annual Review of Economics*, *3*, 513–536.

Romm, A. (2014). Implications of capacity reduction and entry in many-to-one stable matching. *Social Choice and Welfare*, *43*(4), 851–875.

Roth, A. E. (1984). The evolution of the labor market for medical interns and residents: A case study in game theory. *Journal of Political Economy*, *92*(6), 991–1016.

Roth, A. E. (1985). The college admissions problem is not equivalent to the marriage problem. *Journal of Economic Theory*, *36*, 277–288.

Roth, A. E. (2008). Deferred acceptance algorithms: History, theory, practice, and open questions. *International Journal of Game Theory*, *36*(3-4), 537–569.

Roth, A. E. (2013). What have we learned from market design? In N. Vulkan, A. E. Roth, & Z. Neeman (Eds.), *The Handbook of Market Design* (pp. 7–50). Oxford University Press.

Roth, A. E. (2015). *Who Gets What – And Why: The New Economics of Matchmaking and Market Design*. Houghton Mifflin Harcourt.

Roth, A. E., & Peranson, E. (1999). The redesign of the matching market for American physicians: Some engineering aspects of economic design. *American Economic Review*, *89*(4), 748–780.

Roth, A. E., & Sotomayor, M. (1990). *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*. Cambridge University Press.

Sönmez, T. (2013). Bidding for army career specialties: Improving the ROTC branching mechanism. *Journal of Political Economy*, *121*(1), 186–219.

Sönmez, T., & Switzer, T. B. (2013). Matching with (branch-of-choice) contracts at United States Military Academy. *Econometrica*, *81*(2), 451–488.

Wang, F., & Kominers, S. D. (2020). *Respect for improvements via matching with contracts.* (Harvard University Working Paper)

Yenmez, M. B. (2018). A college admissions clearinghouse. *Journal of Economic Theory*, *176*, 859–885.