

Respect for Improvements and Comparative Statics in Matching Markets*

Scott Duke Kominers
Harvard Business School & Department of Economics
Harvard University

July 23, 2020

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.

JEL Classification: C78, D47, D82

Keywords: Matching, Market entry, Respect for improvements

*The author appreciates the helpful comments of Péter Biró, Rachel Bayefsky, Vincent Crawford, Piotr Dworzak, Federico Echenique, John William Hatfield, Emir Kamenica, Shengwu Li, Michael Ostrovsky, Alvin Roth, Larry Samuelson, Tayfun Sönmez, Alexander Teytelboym, Bertan Turhan, Bumin Yenmez, and especially Jiafeng Chen, Ravi Jagadeesan, Shira Li, Jonathan Ma, Franklyn Wang, and other members of the Lab for Economic Design. He gratefully acknowledges 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.

1 Introduction

In their seminal 1962 paper, Gale and Shapley introduced the following 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 and 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 and Peranson (1999)), school choice (Balinski and Sönmez (1999); Abdulkadiroğlu and Sönmez (2003); Pathak (2011)), and the assignment of cadets to branches of military service (Sönmez and Switzer (2013); Sönmez

(2013)).¹

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 and Crawford (1982); Gale and Sotomayor (1985); Roth and 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 and Rothblum (2002); Hatfield and Milgrom (2005); Biró et al. (2008); Ostrovsky (2008); Hatfield and Kominers (2013); Chambers and 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 and Crawford (1982); Crawford (1991); Blum et al. (1997); Blum and Rothblum (2002); Hatfield and Milgrom (2005); Biró et al. (2008); Hatfield and Kominers (2013); Fleiner et al. (2018)) or on broad results characterizing the structure of the set of stable matchings (Gale and Sotomayor (1985); Ostrovsky (2008); Chambers and Yenmez (2017)).

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 and Sönmez (1999)).³ 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.

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

²As 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 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.

³See also the work of Wang and Kominers (2020), which uses the theory of matching with contracts to show a particularly strong link between respect for improvements and strategy-proofness.

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 and Sönmez (2016) introduced to model matching with diversity constraints. In particular, we show that expanding the capacity 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 $\{\text{Alex}, \text{Bob}, \text{Charlie}\}$ and two women $\{\text{Xena}, \text{Yvette}\}$, who rank each other as follows (with the option \emptyset representing being unmatched):

$$\begin{array}{ll}
 \text{Alex} : \text{Yvette} \succ \text{Xena} \succ \emptyset & \text{Xena} : \text{Alex} \succ \text{Bob} \succ \text{Charlie} \succ \emptyset \\
 \text{Bob} : \text{Yvette} \succ \text{Xena} \succ \emptyset & \text{Yvette} : \text{Alex} \succ \text{Bob} \succ \emptyset. \\
 \text{Charlie} : \text{Xena} \succ \text{Yvette} \succ \emptyset &
 \end{array} \tag{1}$$

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 that 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 further that **Zelda**

is especially popular⁴ so that preferences are now given by

$$\begin{array}{ll}
\text{Alex : } \text{Zelda} \succ \text{Yvette} \succ \text{Xena} \succ \emptyset & \text{Xena : } \text{Alex} \succ \text{Bob} \succ \text{Charlie} \succ \emptyset \\
\text{Bob : } \text{Zelda} \succ \text{Yvette} \succ \text{Xena} \succ \emptyset & \text{Yvette : } \text{Alex} \succ \text{Bob} \succ \emptyset \\
\text{Charlie : } \text{Zelda} \succ \text{Xena} \succ \text{Yvette} \succ \emptyset & \text{Zelda : } \text{Alex} \succ \text{Bob} \succ \emptyset.
\end{array} \tag{2}$$

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

$$\begin{array}{ll}
\text{Alex : } \text{Zelda} \succ \text{Yvette} \succ \text{Xena} \succ \emptyset & \text{Xena : } \text{Alex} \succ \text{Bob} \succ \text{Charlie} \succ \emptyset \\
\text{Bob : } \text{Zelda} \succ \text{Yvette} \succ \text{Xena} \succ \emptyset & \text{Yvette : } \text{Alex} \succ \text{Bob} \succ \emptyset \\
\text{Charlie : } \text{Zelda} \succ \text{Xena} \succ \text{Yvette} \succ \emptyset & \text{Zelda : } \emptyset \succ \text{Alex} \succ \text{Bob},
\end{array} \tag{3}$$

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

⁴Perhaps she has an entire videogame franchise to her name.

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 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 and 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 \succsim_i , so that if $j \succsim_i k$ 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 $\succsim_{I'} \equiv (\succsim_i)_{i \in I'}$.

3.1 Stable Matchings

A (*marriage*) *matching* is a map $\mu : I \rightarrow I \cup \{\emptyset\}$ 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 $\mu(i) \succsim_i \emptyset$. 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 and 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 :

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

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 and Crawford (1982); Gale and Sotomayor (1985); Roth and Sotomayor (1990, p. 44)). *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,

$$\tilde{\mu}^*(m) \succsim_m \mu^*(m).^5 \quad (4)$$

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 and 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 and Crawford (1982); Blum et al. (1997)).⁶

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 $\check{\succ}_w$, and the relative rankings of all other men are left unchanged. Formally:

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

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 .

⁵Theorem 1 is also a special case of the comparative static results of Crawford (1991); Hatfield and Milgrom (2005); Ostrovsky (2008); Hatfield and Kominers (2013); Chambers and Yenmez (2017); Yenmez (2018); Fleiner et al. (2018).

⁶Dworczak (2018) tracks similar re-equilibration dynamics in his *Deferred Acceptance with Compensation Chains* algorithms, which generalize deferred acceptance by allowing agents on both sides to make proposals.

Balinski and 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 $\check{\succ}_W$.

Lemma 1 (Balinski and Sönmez (1999)). *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,*

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

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.^{7,8} 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 condition of Definition 1 holds automatically. Thus a (strict) unambiguous improvement for one man can also 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.

4.2 Proof of Theorem 1

As in the theorem statement, 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

$$\check{\succ}_{\tilde{w}} : m_1 \check{\succ}_{\tilde{w}} \cdots \check{\succ}_{\tilde{w}} m_\ell \check{\succ}_{\tilde{w}} \emptyset \check{\succ}_{\tilde{w}} m_{\ell+1} \cdots \quad (5)$$

⁷Imagine the chaos if intentional self-sabotage were a standard feature of the marriage market!

⁸Moreover, respect for unambiguous improvements in some sense uniquely characterizes deferred acceptance (see Theorems 5 and 6 of Balinski and Sönmez (1999)).

and let $\succ_{\tilde{w}}^-$ be the “null” preference relation for \tilde{w} that is consistent with $\tilde{\succ}_{\tilde{w}}$ but treats all men as unacceptable:

$$\succ_{\tilde{w}}^- : \emptyset \succ_{\tilde{w}}^- m_1 \succ_{\tilde{w}}^- \cdots \succ_{\tilde{w}}^- m_\ell \succ_{\tilde{w}}^- m_{\ell+1} \cdots . \quad (6)$$

We let $\bar{\mu}^*$ be the outcome of man-proposing deferred acceptance in the market with agents $I \cup \{\tilde{w}\}$ and preference profile $(\tilde{\succ}_M, \tilde{\succ}_W, \succ_{\tilde{w}}^-)$. 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

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

for each $i \in I$. Indeed, as \tilde{w} treats all men as unacceptable under $\succ_{\tilde{w}}^-$, 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 $\succ_{\tilde{w}}^-$ 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,

$$\tilde{\mu}^*(m) \tilde{\succ}_m \bar{\mu}^*(m) \quad (8)$$

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

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 and Sotomayor (1985). Like us, Gale and Sotomayor (1985) treated \tilde{w} ’s entry as

equivalent to transforming \tilde{w} 's preferences from a null relation to $\tilde{\succ}_{\tilde{w}}$. Unlike in our argument, however, Gale and Sotomayor (1985) derived the comparative static through appeal to a result showing that the set of stable outcomes has on lattice structure (see Roth and 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 and Shapley (1962); Roth (1985); Roth and 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 and Knoer (1981); Kelso and Crawford (1982); Hatfield and Milgrom (2005)). Specifically, we work with the *slot-specific preference* structure introduced by Kominers and 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:

$$s^1 : (i, \$L) \succ (i, \$H) \succ (j, \$L) \succ (j, \$H) \succ \dots$$

$$s^2 : (i, \$L) \succ (j, \$L) \succ \dots$$

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 and Sönmez (2014); Dur et al. (2016, 2018)).

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 $e(x) \in E$ and firm $f(x) \in F$, 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 $e(\cdot)$ and $f(\cdot)$ to sets of contracts $Y \subseteq X$ by setting $e(Y) \equiv \cup_{y \in Y} \{e(y)\}$ and $f(Y) \equiv \cup_{y \in Y} \{f(y)\}$. For $Y \subseteq X$, we denote $Y_e \equiv \{y \in Y : e(y) = e\}$ and $Y_{E'} \equiv \cup_{e \in E'} Y_e$; analogously, we denote $Y_f \equiv \{y \in Y : f(y) = f\}$ 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 $X_e \equiv \{x \in X : e(x) = e\}$ and, as before, \emptyset is an “outside option,” which 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 $X_f \equiv \{x \in X : f(x) = f\}$. Slots $s \in S_f$ have (linear) preference orders \succ_s (with weak orders \succ_s) over contracts in X_f . As with employees, we assume that each slot $s \in S_f$ ranks an outside option \emptyset which 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 $y \in Y_e$).⁹

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* \triangleright^f . We denote $S_f \equiv \{s_f^1, \dots, s_f^{q_f}\}$ with the understanding that $s_f^\ell \triangleright^f s_f^{\ell+1}$ unless otherwise noted. The interpretation of \triangleright^f is that if $s \triangleright^f s'$ 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 $Y \setminus Y_{e(x^1)}$ of contracts in Y with employees other than $e(x^1)$.
- 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 $Y \setminus Y_{e(\{x^1, \dots, x^{\ell-1}\})}$.

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 .

⁹To simplify our exposition and notation, we treat individual contracts as interchangeable with singleton contract sets.

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., $|Y_f| \leq q_f$ for each $f \in F$.

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 $e \in \mathbf{e}(Z)$.

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 and Sönmez (2016) showed that stable outcomes exist under slot-specific preferences, and can be found via the following (*employee-proposing*) *cumulative offer process* (Kelso and Crawford (1982); Hatfield and Milgrom (2005)), which generalizes deferred acceptance:¹⁰

CUMULATIVE OFFER PROCESS

¹⁰The 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. (2018)). For consistency with Kominers and Sönmez (2016), we state the cumulative offer process with a single agent proposing in each step; Hirata and Kasuya (2014) showed 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.

Step 1. Some employee $e^1 \in E$ proposes his or her most-preferred contract, $x^1 \in X_{e^1}$. Firm $f(x^1)$ holds x^1 if $x^1 \in C^{f(x^1)}(\{x^1\})$, and rejects x^1 otherwise. Set $A_{f(x^1)}^2 = \{x^1\}$, and set $A_{f'}^2 = \emptyset$ for each $f' \neq f(x^1)$; 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 his or her most-preferred contract that has not yet been rejected, $x^\ell \in X_{e^\ell}$. Firm $f(x^\ell)$ holds the contracts in $C^{f(x^\ell)}(A_{f(x^\ell)}^\ell \cup \{x^\ell\})$ and rejects all other contracts in $A_{f(x^\ell)}^\ell \cup \{x^\ell\}$; firms $f' \neq f(x^\ell)$ continue to hold all contracts they held at the end of Step $\ell - 1$. Set $A_{f(x^\ell)}^{\ell+1} = A_{f(x^\ell)}^\ell \cup \{x^\ell\}$, and set $A_{f'}^{\ell+1} = A_{f'}^\ell$ for each $f' \neq f(x^\ell)$.

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 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.¹¹

Although stable outcomes exist under slot-specific preferences, the set of stable outcomes under slot-specific preferences does not have lattice structure (see Kominers and Sönmez (2016); Hatfield and Kominers (2019)). Consequently the approach Gale and 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.

¹¹Here 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 and Sönmez (2016), as well as Hirata and Kasuya (2014)).

5.5 Respect for Improvements Under Slot-Specific Preferences

Following Kominers and Sönmez (2016), $\hat{\succ}_S$ 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. 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 $s \in S$:

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

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 his or her cumulative offer process outcome under $(\succ_E, \hat{\succ}_S)$ to that under $(\succ_E, \check{\succ}_S)$.

Lemma 2 (Kominers and Sönmez (2016)). *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 his or her assignment under \hat{Y}^* to that under \check{Y}^* ; that is,*

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

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 his or her assignment under \tilde{Y}^* to his or her assignment under Y^* ; that is,*

$$\tilde{Y}_e^* \succ_e Y_e^*.$$

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 to a firm that previously had none. Theorem 2 also has practical implications: Perhaps most naturally, the result means that under 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 $y \tilde{\succ}_r y'$ 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 $y \tilde{\succ}_s \tilde{x}$ 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 his or her assignment under \tilde{Y}^* to his or her assignment under Y^* ; that is,*

$$\tilde{Y}_e^* \tilde{\succ}_e Y_e^*.$$

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.¹²

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

¹²We cannot in general add contracts at arbitrary positions for multiple agents, however, as the resulting change in preferences might not be an unambiguous improvement.

$y \tilde{\succ}_r y'$ 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 $\mathbf{e}(\tilde{x})$: employee $\mathbf{e}(\tilde{x})$ 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 $\mathbf{e}(\tilde{x})$ (weakly) prefers his or her assignment under \tilde{Y}^* to his or her assignment under Y^* ; that is,*

$$\tilde{Y}_{\mathbf{e}(\tilde{x})}^* \tilde{\succ}_{\mathbf{e}(\tilde{x})} Y_{\mathbf{e}(\tilde{x})}^*.$$

5.7 Discussion

Our use of the respect for improvements result (Lemma 2) here is similar to an early application Kominers and Sönmez (2014) presented for the slot-specific preferences framework. Indeed, Kominers and Sönmez (2014) use 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 and Sönmez (2014) argument consists 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 and Yenmez (2017) showed a result analogous to Theorem 2 for ways of “expanding” choice rules that satisfy a regularity condition called *path-independence* (Aizerman and Malishevski (1981)). Yenmez (2018) extended the Chambers and Yenmez (2017) result still further to cover choice rules that can be modified—or in the language of

Hatfield and Kominers (2019), *completed*—in ways that make them path-independent (see also Echenique and Yenmez (2015); Kamada and Kojima (2015)). Slot-specific preferences are not path-independent, in general, but *do* have path-independent completions (see Hatfield and Kominers (2019)). However, the additions of slots and contracts considered in Theorems 2–4 are *not* expansions in the sense of Chambers and Yenmez (2017) and Yenmez (2018); hence, our results here extend Chambers and 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.¹³

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.¹⁴

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 and Sotomayor (1985); Roth and 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 and Yenmez (2017)). It is unclear whether our respect for improvements-based approach can also be used to prove the refinements just

¹³That 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. (forthcoming)).

¹⁴Likewise, we can extend the entry comparative static to some matching settings under weakened substitutability conditions using the respect for improvements results of Afacan (2017).

described. At minimum, such an argument would seem to require a sharper version of the respect for improvements theorem.

References

- Abdulkadiroğlu, Atila**, “College Admissions with Affirmative Action,” *International Journal of Game Theory*, 2005, *33*, 535–549.
- **and Tayfun Sönmez**, “School Choice: A Mechanism Design Approach,” *American Economic Review*, 2003, *93*, 729–747.
- Afacan, Mustafa Oğuz**, “Some Further Properties of the Cumulative Offer Process,” *Games and Economic Behavior*, 2017, *104*, 656–665.
- Aizerman, Mark A. and Andrew V. Malishevski**, “General theory of best variants choice: Some aspects,” *IEEE Transactions on Automatic Control*, 1981, *26* (5), 1030–1040.
- Balinski, Michel and Tayfun Sönmez**, “A tale of two mechanisms: Student placement,” *Journal of Economic Theory*, 1999, *84*, 73–94.
- Biró, Péter, Katarína Cechlárová, and Tamás Fleiner**, “The dynamics of stable matchings and half-matchings for the stable marriage and roommates problems,” *International Journal of Game Theory*, 2008, *36* (3-4), 333–352.
- Blum, Yosef, Alvin E. Roth, and Uriel G. Rothblum**, “Vacancy chains and equilibration in senior-level labor markets,” *Journal of Economic Theory*, 1997, *76* (2), 362–411.
- **and Uriel G. Rothblum**, ““Timing Is Everything” and Marital Bliss,” *Journal of Economic Theory*, 2002, *103* (2), 429–443.
- Chambers, Christopher P. and M. Bumin Yenmez**, “Choice and matching,” *American Economic Journal: Microeconomics*, 2017, *9* (3), 126–147.

- Crawford, Vincent P.**, “Comparative statics in matching markets,” *Journal of Economic Theory*, 1991, *54* (2), 389–400.
- **and Elsie Marie Knoer**, “Job Matching with Heterogeneous Firms and Workers,” *Econometrica*, 1981, *49* (2), 437–450.
- Dupuy, Arnaud, Alfred Galichon, Sonia Jaffe, and Scott Duke Kominers**, “Taxation in Matching Markets,” *International Economic Review*, forthcoming.
- Dur, Umut, Parag A. Pathak, and Tayfun Sönmez**, “Explicit vs. Statistical Preferential Treatment in Affirmative Action: Theory and Evidence from Chicago’s Exam Schools,” 2016.
- **, Scott Duke Kominers, Parag A. Pathak, and Tayfun Sönmez**, “Reserve Design: Unintended Consequences and The Demise of Boston’s Walk Zones,” *Journal of Political Economy*, 2018, *126* (6), 2457–2479.
- Dworzak, Piotr**, “Deferred Acceptance with Compensation Chains,” 2018. Northwestern University Working Paper.
- Echenique, Federico and M. Bumin Yenmez**, “How to control controlled school choice,” *American Economic Review*, 2015, *105*, 2679–2694.
- Fleiner, Tamás, Zsuzsanna Jankó, Akihisa Tamura, and Alexander Teytelboym**, “Trading networks with bilateral contracts,” 2018. Oxford University Working Paper.
- Gale, David and Lloyd S. Shapley**, “College Admissions and the Stability of Marriage,” *American Mathematical Monthly*, 1962, *69* (1), 9–15.
- **and Marilda A. O. Sotomayor**, “Some remarks on the stable matching problem,” *Discrete Applied Mathematics*, 1985, *11* (3), 223–232.
- Hafalir, Isa Emin, M. Bumin Yenmez, and Muhammed Ali Yildirim**, “Effective Affirmative Action in School Choice,” *Theoretical Economics*, 2013, *8*, 325–363.

- Hatfield, John William and Paul Milgrom**, “Matching with Contracts,” *American Economic Review*, 2005, *95* (4), 913–935.
- **and Scott Duke Kominers**, “Vacancies in supply chain networks,” *Economics Letters*, 2013, *119* (3), 354–357.
- **and –**, “Hidden Substitutes,” 2019. Harvard University Working Paper.
- , – , **and Alexander Westkamp**, “Stability, Strategy-Proofness, and Cumulative Offer Mechanisms,” 2018. Harvard University Working Paper.
- Hirata, Daisuke and Yusuke Kasuya**, “Cumulative offer process is order-independent,” *Economics Letters*, 2014, *124* (1), 37–40.
- Jr., Alexander S. Kelso and Vincent P. Crawford**, “Job Matching, Coalition Formation, and Gross Substitutes,” *Econometrica*, 1982, *50* (6), 1483–1504.
- Kamada, Yuichiro and Fuhito Kojima**, “Efficient Matching Under Distributional Constraints: Theory and Applications,” *American Economic Review*, 2015, *105*, 67–99.
- Kojima, Fuhito**, “School Choice: Impossibilities for Affirmative Action,” *Games and Economic Behavior*, 2012, *75*, 685–693.
- Kominers, Scott Duke, Alexander Teytelboym, and Vincent P. Crawford**, “An Invitation to Market Design,” *Oxford Review of Economic Policy*, 2017, *33* (4), 541–571.
- **and Tayfun Sönmez**, “Designing for Diversity in Matching,” 2014. Harvard University Working Paper.
- **and –**, “Matching with Slot-Specific Priorities: Theory,” *Theoretical Economics*, 2016, *11* (2), 683–710.
- Ostrovsky, Michael**, “Stability in Supply Chain Networks,” *American Economic Review*, 2008, *98* (3), 897–923.

- Pathak, Parag A.**, “The mechanism design approach to student assignment,” *Annual Review of Economics*, 2011, *3*, 513–536.
- Romm, Assaf**, “Implications of capacity reduction and entry in many-to-one stable matching,” *Social Choice and Welfare*, 2014, *43* (4), 851–875.
- Roth, Alvin E.**, “The Evolution of the Labor Market for Medical Interns and Residents: A Case Study in Game Theory,” *Journal of Political Economy*, 1984, *92* (6), 991–1016.
- , “The College Admissions Problem is not Equivalent to the Marriage Problem,” *Journal of Economic Theory*, 1985, *36*, 277–288.
- , “Deferred Acceptance Algorithms: History, Theory, Practice, and Open Questions,” *International Journal of Game Theory*, 2008, *36* (3-4), 537–569.
- , “What Have We Learned From Market Design?,” in Nir Vulkan, Alvin E. Roth, and Zvika Neeman, eds., *The Handbook of Market Design*, Oxford University Press, 2013, pp. 7–50.
- , *Who Gets What – And Why: The New Economics of Matchmaking and Market Design*, Houghton Mifflin Harcourt, 2015.
- **and Elliott Peranson**, “The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design,” *American Economic Review*, 1999, *89* (4), 748–780.
- **and Marilda Sotomayor**, *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*, Vol. 18 of *Econometric Society Monographs*, Cambridge University Press, 1990.
- Sönmez, Tayfun**, “Bidding for Army Career Specialties: Improving the ROTC Branching Mechanism,” *Journal of Political Economy*, 2013, *121*, 186–219.
- **and Tobias B. Switzer**, “Matching with (Branch-of-Choice) Contracts at United States Military Academy,” *Econometrica*, 2013, *81*, 451–488.

Wang, Franklyn and Scott Duke Kominers, “Respect for Improvements via Matching with Contracts,” 2020. Harvard University Working Paper.

Yenmez, M. Bumin, “A college admissions clearinghouse,” *Journal of Economic Theory*, 2018, *176*, 859–885.