

Designing for Diversity: Matching with Slot-Specific Priorities*

Scott Duke Kominers
Becker Friedman Institute
University of Chicago

Tayfun Sönmez
Department of Economics
Boston College

July 2, 2012

Abstract

To encourage diversity, branches may vary contracts' priorities across slots. The agents who match to branches, however, have preferences only over match partners and contractual terms. *Ad hoc* approaches to resolving agents' indifferences across slots in the Chicago and Boston school choice programs have introduced biases, which can be corrected with more careful market design.

Slot-specific priorities can fail the substitutability condition typically crucial for outcome stability. Nevertheless, an embedding into a one-to-one agent–slot matching market shows that stable outcomes exist and can be found by a cumulative offer mechanism that is strategy-proof and respects unambiguous improvements in priority.

JEL classification: C78, D63, D78.

Keywords: Market Design, Matching with Contracts, Stability, Strategy-Proofness, School Choice, Affirmative Action.

*The authors are deeply grateful to Parag Pathak for his comments and for providing access to data from the Chicago public school choice program. They appreciate the helpful suggestions of Eric Budish, John William Hatfield, Ellen Kominers, Sonia Jaffe, Jonathan Levin, Alexandru Nichifor, Philip J. Reny, Assaf Romm, Alvin E. Roth, Alexander Westkamp, seminar audiences at Chicago, Haifa, and St. Andrews, and members of the Measuring and Interpreting Inequality Working Group. Kominers gratefully acknowledges the support of an AMS–Simons travel grant and the Human Capital and Economic Opportunity Working Group sponsored by the Institute for New Economic Thinking.

1 Introduction

The *deferred acceptance algorithm* of Gale and Shapley (1962) has been applied in a variety of real-world matching problems, stabilizing labor markets and school choice programs, and contributing significantly to the emergence and success of the field of *market design*. Until the last decade, deferred acceptance was primarily applied in two-sided, many-to-one matching markets—such as the United States doctor–residency match (Roth (1984); Roth and Peranson (1999))—in which both sides of the market may behave strategically. This trend has changed since the introduction of the *student placement* (Balinski and Sönmez (1999)) and *school choice* (Abdulkadiroğlu and Sönmez (2003)) models, which, following school matching traditions in the United States and elsewhere, focus on settings in which seats at institutions are simply a set of objects to be allocated to agents and institution has an exogenous *priority* ranking over students which determines claims over seats.

As Dubins and Freedman (1981) and Roth (1982) showed, truth-telling is a dominant strategy for all agents when *agent-proposing* deferred acceptance is run. Meanwhile, there exists no stable mechanism that is dominant-strategy incentive compatible for institutions (Roth (1982)). Market designers have thus advocated the use of agent-proposing deferred acceptance, even though use of this mechanism may conflict with the interests of institutions.¹

The case for agent-proposing deferred acceptance is in a sense stronger for priority matching applications than it is for two-sided labor markets, as in such settings agent-proposing deferred acceptance is dominant-strategy incentive compatible for all strategic players, i.e. it is *strategy-proof*.² Stability and strategy-proofness are attractive in priority matching: the former has a natural normative interpretation as *elimination of justified envy*, and the latter

¹In the framework of Gale and Shapley (1962), the proposing side of the market is significantly advantaged: the deferred acceptance algorithm produces an outcome which the proposing side of the market weakly prefers to all other stable outcomes; this outcome is also pessimal among stable outcomes for the non-proposing side of the market.

²As before, conflict of interest between agents and institutions arises, but there is often less concern for institution welfare in priority matching contexts. Institution welfare is not irrelevant, however, as it can affect broader policy goals: Hatfield et al. (2012) show, for example, how welfare effects of school choice mechanism selection can impact schools' incentives for self-improvement.

levels the playing field by eliminating gains to strategic sophistication and enables the collection of true preference data for planning purposes (Abdulkadiroğlu et al. (2006); Pathak and Sönmez (2008)). Because of these advantages, a large number of school districts have recently adopted centralized school choice programs based on the agent-proposing deferred acceptance.³ Many of these school districts (e.g., Boston, New York City, and Chicago) are concerned with issues of student diversity and have therefore embedded affirmative action systems into their school choice programs. At present, however, these diversity efforts are handled in a somewhat *ad hoc* manner.

In this paper, we observe that diversity, financial aid, or other concerns might cause priorities to vary across slots at a given institution. We argue that to effectively handle these *slot-specific priority* structures, market designers should go beyond the traditional deferred acceptance algorithm and use a more detailed design approach based on the Kelso and Crawford (1982)/Hatfield and Milgrom (2005) theory of *many-to-one matching with contracts*.

To make this case, we first illustrate how existing *ad hoc* deferred acceptance implementations impact student welfare in the Boston and Chicago school choice programs. Then, we introduce a model of matching with slot-specific priorities, which embeds classical priority matching frameworks (e.g., Balinski and Sönmez (1999); Abdulkadiroğlu and Sönmez (2003)), models of affirmative action (e.g., Kojima (2012); Hafalir et al. (forthcoming)), and the cadet–branch matching framework (Sönmez and Switzer (2011); Sönmez (2011)).

We advocate for a specific implementation of the *cumulative offer mechanism* of Hatfield and Milgrom (2005) and Hatfield and Kojima (2010), which generalizes agent-proposing deferred acceptance. Previous priority matching models have relied on the existence of agent-optimal stable outcomes to guarantee that this mechanism is strategy-proof. In markets

³These school choice reforms include assignment of high school students in New York City in 2003 Abdulkadiroğlu et al. (2005b, 2009), assignment of K–12 students to public schools in Boston in 2005 Abdulkadiroğlu et al. (2005a), assignment of high school students to selective enrollment schools in Chicago in 2009 Pathak and Sönmez (forthcoming), and assignment of K–12 students to public schools in Denver in 2012. Perhaps most significantly, a version of the agent-proposing deferred acceptance has been recently adopted by all (more than 150) local authorities in England (Pathak and Sönmez (forthcoming)).

with slot-specific priorities, however, agent-optimal stable outcomes may not exist. Nevertheless, as we show, the cumulative offer mechanism is still strategy-proof in such markets; this observation is perhaps the most important theoretical contribution of our work. We show moreover that the cumulative offer mechanism has two other features essential for applications: the cumulative offer mechanism yields stable outcomes and respects unambiguous improvements of agent priority.⁴

Our work demonstrates that the existence of a plausible mechanism for real-world many-to-one matching with contracts does not rely on the existence of an agent-optimal stable outcome. However, it is worth emphasizing that the cumulative offer mechanism selects the agent-optimal stable outcome whenever such an outcome exists. This is important because the existence of an agent-optimal stable outcome can be considered quite relevant for applications because it removes all conflict of interest among agents in the context of stable assignment. The existence of an agent-optimal stable outcome in our general model may depend on several factors, including the number of different contractual arrangements each agent can have with each institution and the *precedence order* according to which institutions prioritize individual slots above others.⁵

Finally, we note that our paper also has a methodological contribution: In general, slot-specific priorities fail the *substitutability* condition that has so far been key in analysis of most two-sided matching with contracts models (Kelso and Crawford (1982); Hatfield and Milgrom (2005); see also Adachi (2000); Fleiner (2003); Echenique and Oviedo (2004)).⁶ Moreover, slot-specific priorities may fail the *unilateral substitutability* condition of Hatfield and Kojima (2010) that has been central to the analysis of cadet–branch matching (Sönmez and Switzer (2011); Sönmez (2011)). Nevertheless, the priority structure in our model gives rise to

⁴These conclusions allow us to re-derive several main results of the innovative Hafalir et al. (forthcoming) approach to welfare-enhancing affirmative-action.

⁵As we show in an application of our model (Proposition 5), the United States Military Academy cadet–branch matching mechanism in a sense uses the unique precedence order under which the cumulative offer mechanism is agent-optimal.

⁶Thus, in particular, our model falls outside of the domain which Echenique (2012) has shown can be handled with only the Kelso and Crawford (1982) *matching with salaries* framework (see also Kominers (2012)).

a naturally associated *one-to-one* model of agent–slot matching (with contracts). As the agent–slot matching market is one-to-one, it trivially satisfies the *substitutability* condition. It follows that the set of outcomes stable in the agent–slot market (called *slot-stable* outcomes to avoid confusion) has an agent-optimal element. We show that each slot-stable outcome corresponds to a stable outcome⁷; moreover, we show that the cumulative offer mechanism in the “true” matching market gives the outcome which corresponds to the agent-optimal slot-stable outcome in the agent–slot matching market. We utilize these relationships heavily in proving our main results.

The remainder of this paper is organized as follows. In Section 2, we discuss the Chicago and Boston school choice programs, illustrating the importance of careful design in markets with slot-specific priorities. We present our model of matching with slot-specific priorities in Section 3. Then, in Section 4, we introduce the agent–slot matching market and derive key properties of the cumulative offer mechanism. In Section 5, we revisit the affirmative-action applications and briefly discuss applications to cadet–branch matching. Most proofs are contained in the Appendix.

2 Motivating Applications

First advocated by Friedman (1955, 1962), *school choice* programs aim to enable parents to choose which schools their children attend. There is significant tension between the proponents of school choice and the proponents of alternative, *neighborhood assignment* systems based on students’ home addresses.

The mechanics of producing the assignment of students to school seats received little attention in school choice debates until Abdulkadiroğlu and Sönmez (2003) showed important shortcomings of several mechanisms adopted by United States school districts. Of particular concern was the vulnerability of school choice to preference manipulation: while parents were

⁷The converse result is not true, in general—there may be stable outcomes which are not associated to slot-stable outcomes.

allowed to express their preferences on paper, they were implicitly forced to play sophisticated admission games. Once this flaw became clear, several school districts adopted the (student-proposing) deferred acceptance mechanism, which was invented by Gale and Shapley (1962) and proposed as a school choice mechanism by Abdulkadiroğlu and Sönmez (2003).⁸

Deferred acceptance-based mechanisms have been successful in part because they are fully flexible regarding the choice of student priorities at schools. In particular, priority rankings may vary across schools; hence, students can be given the option of school choice while retaining some priority for their neighborhood schools. Thus, deferred acceptance mechanisms provide a natural opportunity for policymakers to balance the concerns of both school choice and neighborhood assignment advocates.

However, deferred acceptance-style mechanisms are designed under the assumption that student priorities are identical across a given school's seats. While this assumption is natural for some school choice applications, it fails in several important cases: admissions to selective high schools in Chicago; K–12 school admissions in Boston; and public high school admission in New York. We next describe

- how the Chicago and Boston matching problems differ from the original school choice model of Abdulkadiroğlu and Sönmez (2003),
- how policymakers in Chicago and Boston sought to transform their problems into direct applications of the Abdulkadiroğlu and Sönmez (2003) framework, and
- how these transformations introduced significant, yet hidden biases in the underlying priority structures.

These observations motivate the more general matching model with slot-specific priorities that we introduce in Section 3.

⁸This mechanism is often called the *student-optimal stable matching mechanism*.

2.1 Affirmative Action at Chicago’s Selective High Schools

Since 2009, Chicago Public Schools (CPS) has adopted an affirmative action plan based on socio-economic status (SES). Although CPS initially adopted a version of the Boston mechanism for selective enrollment high school admissions, they immediately abandoned it in favor of a deferred acceptance-based approach.⁹ Under the new assignment plan, the SES of each student is determined based on home address; students are then divided into four roughly evenly sized tiers:¹⁰ In 2009,

- 135,716 students living in 210 Tier 1 (lowest-SES) tracts had a median family income of \$30,791,
- 136,073 students living in 203 Tier 2 tracts had a median family income of \$41,038,
- 136,378 students living in 226 Tier 3 tracts had a median family income of \$54,232, and
- 136,275 students living in 235 Tier 4 (highest-SES) tracts had a median family income of \$76,829.

Students in Chicago who apply to selective enrollment high schools take an admissions test as part of their application, and this test is used to determine a *composite score*.¹¹ Students from high-SES tiers typically have higher composite scores, and CPS has set aside seats as reserved for low-SES students in order to prevent the elite schools from becoming accessible only to children from wealthy neighborhoods. In order to implement this objective, CPS adopted the following priority structure at each of the nine selective enrollment high schools:

- priority for 40% of the seats is determined by students’ composite scores, while

⁹Pathak and Sönmez (forthcoming) present a detailed account of this midstream reform.

¹⁰SES scores are uniform across census tracts, and are based on median family income, average adult educational attainment, percentage of single-parent households, percentage of owner-occupied homes, and percentage of non-English speakers.

¹¹If two students have the same score, then the younger student is coded by CPS as having a higher composite score.

- 15% of the seats are set aside for students in each of the four SES tiers \mathbf{t} , with composite score determining relative priority among students in Tier \mathbf{t} .¹²

Because these priorities over seats are not uniform within schools, the Abdulkadiroğlu and Sönmez (2003) school choice model does not fully capture all aspects of the Chicago admissions problem. In order to implement deferred acceptance despite this difficulty, the CPS matching algorithm treats each selective enrollment high school as five hypothetical schools: The set S_b of seats at each school b is partitioned into subsets

$$S_b = S_b^o \cup S_b^4 \cup S_b^3 \cup S_b^2 \cup S_b^1.$$

The seats in S_b^o are “open seats,” for which students’ priorities are determined entirely by composite scores. Seats in S_b^t are “reserved” for students of Tier \mathbf{t} —they give Tier \mathbf{t} students priority over other students, and use composite scores to rank students within Tier \mathbf{t} . Each set of seats is viewed as a separate “school” within the CPS algorithm. Because seat priorities are uniform within each set S_b^e , the set of hypothetical schools satisfies the Abdulkadiroğlu and Sönmez (2003) requirement of uniform within-school priorities.

However, CPS students i submit preferences P^i over *schools*, and are indifferent among seats at a given school $b \in B$. That is, if we denote a contract representing that i holds a seat in S_b^e by $\langle i; s_b^e \rangle$, then any Tier \mathbf{t} student i is indifferent between

$$\langle i; s_b^o \rangle, \quad \langle i; s_b^4 \rangle, \quad \langle i; s_b^3 \rangle, \quad \langle i; s_b^2 \rangle, \quad \text{and} \quad \langle i; s_b^1 \rangle,$$

and prefers all these these contracts to any contract of the form $\langle i; s_{b'}^e \rangle$ if (and only if) i prefers school b to school b' (i.e. $bP^i b'$). As the Abdulkadiroğlu and Sönmez (2003) model requires that students have *strict* preferences over schools, the CPS matching algorithm must

¹²The priority structure we describe here was used in the 2010–2011 CPS match. The CPS priority structure has been revised slightly for school year 2012-2013. Then, 5% of the seats are reserved for a second round to be hand picked by principals, the fraction of open competition seats is reduced to 28.5%, and the fraction of reserved seats is increased to 16.625% for each SES tier.

convert students’ true preferences P^i into strict, “extended” preferences \tilde{P}^i over the full set of hypothetical schools. In practice, CPS does this by assuming that

$$\langle i; s_b^o \rangle \tilde{P}^i \langle i; s_b^4 \rangle \tilde{P}^i \langle i; s_b^3 \rangle \tilde{P}^i \langle i; s_b^2 \rangle \tilde{P}^i \langle i; s_b^1 \rangle$$

for each student i and school b .¹³ That is, CPS assumes that students most prefer open seats, and then rank reserved seats. CPS thus picks for each student the (unique) preference ranking consistent with the student’s submitted preferences in which the open seats at each school are ranked immediately above the reserved seats. With this transformation of preferences, the Chicago problem finally fits within the standard school choice framework of Abdulkadiroğlu and Sönmez (2003).

Being an affirmative action plan, the priority structure in Chicago is designed so that students of lower SES tiers receive favorable treatment. What may be less clear is that the transformation used in the CPS implementation of deferred acceptance provides additional advantages to lower SES tier students.

To understand this bias, we consider a simple example in which there is only one school. The CPS matching algorithm first assigns the open seats and subsequently assigns the reserved seats. Therefore (under the 2010-2011 CPS plan), 40% of the seats are assigned to the students—from any SES tier—with highest composite scores, and the remaining 60% of seats are shared evenly among the four SES tiers. Hence, students in each tier have access to 15% of the seats plus some fraction of the open seats depending on their composite scores. In order to see how this treatment favors for lower tiers, we consider an alternative mechanism in which reserved seats are allocated before open competition seats. Under this counterfactual mechanism:

First, the reserved seats are allocated to the students in each SES tier with the highest composite scores. Then, the students remaining unassigned are ranked

¹³Note that as seats in S_b^t are reserved for Tier t students, the only “relevant” part of this construction is the fact that $\langle i; s_b^o \rangle \tilde{P}^i \langle i; s_b^t \rangle$ for any Tier t student i .

Tier 4
Tiers 4 and 3
Tiers 4, 3, and 2
Tiers 4, 3, 2, and 1

Figure 1: Top-scoring students in the truncated distribution come from Tier 4; the next-highest score block consists only of students from Tiers 4 and 3, and so forth.

according to their composite scores and admitted in descending order of score until all the open seats are filled.

High-SES students' composite scores dominate low-SES students' scores throughout the relevant part of the score distribution—the number of Tier 4 students with score σ high enough to gain admission is larger than the number of Tier 3 students with score σ , and so forth. In practice, this means that after the reserved seats are filled, high-SES students fill most (if not all) of the open seats. Indeed, once the highest-scoring students in each tier are removed, the score distribution of students vying for the last 40% of seats takes the block form illustrated in Figure 1. Thus, the open seats first fill only with Tier 4 students, then fill with both Tiers 4 and 3, and then fill with students from Tiers 4, 3, and 2. Only after that (if seats remain) do Tier 1 students gain access. Of course, the size of the blocks—and hence the size of the effect of switching to the counterfactual mechanism—is an empirical question. Using actual data from 2010-2011 Chicago school choice admissions program, we now show that our intuition is accurate and that the magnitude of this effect is substantial.

We compare the outcome of the Chicago match under two alternative scenarios:

1. First we consider the current system, in which all open seats are allocated before reserved seats.
2. Second we consider the counterfactual in which all open seats are allocated after reserved seats.

	Tier 4	Tier 3	Tier 2	Tier 1
Prefer Actual	62	50	108	175
Indifferent	3748	4287	4092	3474
Prefer Counterfactual	225	108	43	0

Table 1: Comparison of individual student outcomes.

Of the 4270 seats allocated at nine selective high schools, 771 of them change hands between the two scenarios—what CPS most likely considers a minor coding decision in fact impacts 18% of the seats at Chicago’s selective high schools.

Table 1 shows that Tier 1 students unambiguously prefer the current mechanism; this confirms that the current implementation particularly benefits low-tier students. The effect is not uniform for all members of higher SES tiers. Nevertheless, it is consistent with our simple example in terms of aggregate distribution:

- Of the 151 Tier 2 students who are affected, 108 prefer the current treatment and 43 prefer the counterfactual.
- Of the 158 Tier 3 students who are affected, 50 prefer the current treatment and 108 prefer the counterfactual.
- And of the 287 Tier 4 students who are affected, 62 prefer the current treatment and 225 prefer the counterfactual.

Table 2 shows how the composition of admitted classes varies between the current and counterfactual mechanisms. We observe that Tier 1 students benefit from the current implementation at the expense of students in all three higher tiers. Perhaps surprisingly, most of the benefit appears to be at the expense of Tier 3 students.¹⁴

Finally, Table 3 compares the number of students receiving seats under the counterfactual mechanism to the sizes of schools’ reserved seat blocks. We see that switching to the counterfactual mechanism can in effect convert reserves into quotas: Consistent with the intuition of our informal, one-school example, Tier 1 students receive no open seats under the

¹⁴One reason for this might be the availability of high quality outside options for some members of Tier 4.

	Current Mechanism (Open Slots First)				Effect of Switching (to fill Open Slots Last)			
	Tier 4	Tier 3	Tier 2	Tier 1	Tier 4	Tier 3	Tier 2	Tier 1
1	105	71	47	43	30	-18	-8	-4
2	95	114	70	60	0	24	-14	-10
3	87	78	86	73	-36	16	35	-15
4	106	93	80	68	-21	9	28	-16
5	210	100	78	78	20	-2	-9	-9
6	121	69	45	49	25	-15	-3	-7
7	655	412	291	272	29	37	-38	-28
8	90	47	36	36	3	7	-5	-5
9	92	129	90	94	-27	5	19	3
TOT	1561	1113	823	773	23	63	5	-91

Table 2: Effect of mechanism change on student body composition at each of the nine selective high schools.

	Counterfactual Mechanism (Open Slots Last)				Reserve/Tier
	Tier 4	Tier 3	Tier 2	Tier 1	
1	135	53	39	39	39
2	95	138	56	50	50
3	51	94	121	58	48
4	85	102	108	52	52
5	230	98	69	69	69
6	146	54	42	42	42
7	684	449	253	244	244
8	93	54	31	31	31
9	65	134	109	97	60

Table 3: Seat allocations under the counterfactual mechanism, in comparison to the 15% reserve.

counterfactual mechanism at seven of the nine selective high schools. And again consistent with our example, we see that in four of the nine schools, even Tier 2 students receive no open seats under the counterfactual.

2.2 K–12 Admissions in Boston Public Schools

In the Boston school choice program, the priority of a student for a given school depends on

1. whether the student has a sibling at that school (**sibling priority**),
2. whether the student lives within the (objectively determined) walk-zone of that school (**walk-zone priority**), and
3. a random number used to break ties.

Students with walk-zone priority alone have higher claims for 50% of the seats at their neighborhood schools, but have no priority advantage at other seats. This choice of priority structure in Boston is not arbitrary—it represents a delicate balance between the interests of the school choice and neighborhood assignment advocates.

As in the case of Chicago, Boston school choice priorities are not uniform across schools’ seats. And similarly to the Chicago mechanism, the BPS school choice algorithm treats each school as two hypothetical schools (of half the true capacity), one with walk-zone priority and one without. To convert student preferences over true schools into preferences over hypothetical schools, BPS chooses the (unique) ranking consistent with the ranking of the original schools such that seats with walk-zone priority are ranked above seats without walk-zone priority. Because of this implementation decision, the BPS school choice program in fact systematically favors proponents of school choice over proponents of neighborhood assignment.¹⁵ This observation becomes more interesting in light of the recent January 17, 2012 statement by Boston Mayor Thomas M. Menino:

“I’m committing tonight that one year from now Boston will have adopted a radically different student assignment plan, one that puts a priority on children attending schools closer to their homes.”¹⁶

The current BPS mechanism design appears to be in conflict with Menino’s statement. This conflict can be corrected by reversing the assumed order of the two types of seats, or by adopting a more balanced implementation based on the model we develop in this paper.

2.3 Discussion

The current Chicago school choice mechanism’s treatment of low-SES students might well be consistent with the policy objectives of CPS. However, it is not clear that the current BPS

¹⁵Intuition for this fact is much the same as for the case of Chicago: Under the current system, walk-zone students with high enough priority to receive non-walk-zone seats are systematically awarded walk-zone seats, effectively wasting their high priority draws.

¹⁶See transcript at http://articles.boston.com/2012-01-18/metro/30634585_1_boston-police-officers-city-speech-unemployment-rate.

school choice mechanism is consistent with Mayor Menino’s stated goal of increasing the emphasis on neighborhood assignment. In any event, it is clear that *ad hoc* implementations of deferred acceptance have introduced implicit biases in both the Chicago and Boston school choice programs.

The general model we provide in the remainder of this paper encapsulates the Chicago and Boston school choice settings, and provides machinery which can virtually eliminate the biases we have illustrated.

3 Basic Model

3.1 Agents, Branches, Contracts, and Slots

A matching problem with **slot-specific priorities** is comprised of a set of **agents** I , a set of **branches** B , and a (finite) set of **contracts** X .¹⁷ Each contract $x \in X$ is between an agent $i(x) \in I$ and branch $b(x) \in B$ ¹⁸; we extend the notations $i(\cdot)$ and $b(\cdot)$ to sets of contracts by setting $i(Y) \equiv \cup_{y \in Y} \{i(y)\}$ and $b(Y) \equiv \cup_{y \in Y} \{b(y)\}$.

Each agent $i \in I$ has a (linear) preference order P^i (with weak order R^i) over contracts in $X_i \equiv \{x \in X : i(x) = i\}$. For ease of notation, we assume that each i also ranks a “null contract” \emptyset_i which represents remaining unmatched (and hence is always available), so that we may assume that i ranks *all* the contracts in X_i . We say that the contracts $x \in X$ for which $\emptyset_i P^i x$ are **unacceptable to i** .

Each branch $b \in B$ has a set S_b of slots which can be assigned contracts in $X_b \equiv \{x \in X : b(x) = b\}$. Each slot $s \in S_b$ has a (linear) priority order Π^s (with weak order Γ^s) over contracts in X_b ; for convenience, we use the convention that $X_s \equiv X_b$ for $s \in S_b$. As with agents, we assume that each slot s ranks a “null contract” \emptyset_s which represents remaining

¹⁷Here, a *branch* could represent a branch of the military (as in cadet–branch matching), or a school (as in the Chicago and Boston examples).

¹⁸A contract may have additional “terms” in addition to an agent and a branch. For concreteness, X may be considered a subset of $I \times B \times T$ for some set T of potential contract terms.

unassigned. We set $S \equiv \cup_{b \in B} S_b$.

3.2 Choice and the Order of Precedence

We assume that for each $b \in B$, the slots in S_b are ordered according to a (linear) **order of precedence** \triangleright^b . We denote $S_b \equiv \{s_b^1, \dots, s_b^{q_b}\}$ with $q_b \equiv |S_b|$ and the understanding that $s_b^\ell \triangleright^b s_b^{\ell+1}$ unless otherwise noted. The interpretation of \triangleright^b , defined formally in the choice function construction below, is that if $s_b^\ell \triangleright^b s_b^{\ell'}$ then—whenever possible—branch b fills slot s_b^ℓ before filling $s_b^{\ell'}$.

To simplify our exposition and notation in the sequel, we treat linear orders over contracts as interchangeable with orders over singleton contract sets.

We denote by $\max_{\bar{P}^i} Y$ the \bar{P}^i -maximal element of Y_i , using the convention that $\max_{\bar{P}^i} Y = \emptyset_i$ if $\emptyset_i \bar{P}^i y$ for all $y \in Y_i$; similarly, we denote by $\max_{\bar{\Pi}^s} Y$ the $\bar{\Pi}^s$ -maximal element of Y_s , using the convention that $\max_{\bar{\Pi}^s} Y = \emptyset_s$ if $\emptyset_s \bar{\Pi}^s y$ for all $y \in Y_s$.¹⁹

Agents have **unit demand**, that is, they choose at most one contract from a set of contract offers. We assume also that agents always choose the maximal available contract, so that the choice $C^i(Y)$ of an agent $i \in I$ from contract set $Y \subseteq X$ is defined by

$$C^i(Y) \equiv \max_{P^i} Y.$$

Meanwhile, branches $b \in B$ may be assigned as many as q_b contracts from an offer set $Y \subseteq X$ —one for each slot in S_b —but may hold no more than one contract with a given agent. Slots at branch b are filled in the order of precedence \triangleright^b :

- First, slot s_b^1 is assigned the contract x^1 which is $\Pi^{s_b^1}$ -maximal among contracts in Y .
- Then, slot s_b^2 is assigned the contract x^2 which is $\Pi^{s_b^2}$ -maximal among contracts in the set $Y \setminus Y_{i(x^1)}$ of contracts in Y with agents other than $i(x^1)$.

¹⁹Here, we use the notations \bar{P}^i and $\bar{\Pi}^s$ because we will sometimes need to maximize over orders other than P^i and Π^s .

- This process continues, with each slot s_b^ℓ being assigned the contract x^ℓ which is $\Pi^{s_b^\ell}$ -maximal among contracts in the set $Y \setminus Y_{i(\{x^1, \dots, x^{\ell-1}\})}$.

Formally, the choice (set) $C^b(Y)$ of a branch $b \in B$ from Y is defined by the following algorithm:

1. Let $H_b^0 \equiv \emptyset$, and let $V_b^1 \equiv Y$.
2. For each $\ell = 1, \dots, q_b$:
 - (a) Let $x^\ell \equiv \max_{\Pi^{s_b^\ell}} V_b^\ell$ be the $\Pi^{s_b^\ell}$ -maximal contract in V_b^ℓ .
 - (b) Set $H_b^\ell = H_b^{\ell-1} \cup \{x^\ell\}$ and set $V_b^{\ell+1} = V_b^\ell \setminus Y_{i(x^\ell)}$.
3. Set $C^b(Y) = H_b^{q_b}$.

We say that a contract $x \in Y$ is **assigned to slot** $s_b^\ell \in S_b$ in the computation of $C^b(Y)$ if $\{x\} = H_b^\ell \setminus H_b^{\ell-1}$ in the running of the algorithm defining $C^b(Y)$. If no contract $x \in Y$ is assigned to slot $s_b^\ell \in S_b$ in the computation of $C^b(Y)$, then we say that s_b^ℓ is assigned the null contract $\emptyset_{s_b^\ell}$.

3.3 Stability

An **outcome** is a set of contracts $Y \subseteq X$. We follow the Gale and Shapley (1962) tradition in focusing on match outcomes which are *stable* in the sense that

- neither agents nor branches wish to unilaterally walk away from their assignments, and
- agents and branches cannot benefit by recontracting outside of the match.

Formally, we say that an outcome Y is **stable** if it is

1. **individually rational**— $C^i(Y) = Y_i$ for all $i \in I$ and $C^b(Y) = Y_b$ for all $b \in B$ —and
2. **unblocked**—there does not exist a branch $b \in B$ and **blocking set** $Z \neq C^b(Y)$ such that $Z = C^b(Y \cup Z)$ and $Z_i = C^i(Y \cup Z)$ for all $i \in i(Z)$.

3.4 Conditions on the Structure of Branch Choice

3.4.1 Substitutability Conditions

Definition. A choice function C^b is **substitutable** if for all $z, z' \in X$ and $Y \subseteq X$,

$$z \notin C^b(Y \cup \{z\}) \implies z \notin C^b(Y \cup \{z, z'\}).$$

Hatfield and Milgrom (2005) introduced this substitutability condition, which generalizes the earlier *gross substitutes* condition of Kelso and Crawford (1982). Hatfield and Milgrom (2005) also showed that substitutability is sufficient to guarantee the existence of stable outcomes.²⁰

Choice function substitutability is necessary for the guaranteed existence of stable outcomes in a variety of settings, including many-to-many matching with contracts (Hatfield and Kominers (2010)) and the Ostrovsky (2008) supply chain matching framework (Hatfield and Kominers (2012)). However, substitutability is *not* necessary for the guaranteed existence of stable outcomes in settings where agents have unit demand (Hatfield and Kojima (2008, 2010)). Indeed, as Hatfield and Kojima (2010) showed, the following condition weaker than substitutability suffices not only for the existence of stable outcomes, but also to guarantee that there is no conflict of interest among agents.²¹

Definition. A choice function C^b is **unilaterally substitutable** if

$$z \notin C^b(Y \cup \{z\}) \implies z \notin C^b(Y \cup \{z, z'\})$$

²⁰Their analysis of Hatfield and Milgrom (2005) implicitly assumes **irrelevance of rejected contracts**, the requirement that

$$z \notin C^b(Y \cup \{z\}) \implies C^b(Y) = C^b(Y \cup \{z\})$$

for all $b \in B$, $Y \subseteq X$, and $z \in X \setminus Y$ (Aygün and Sönmez (2012b)). This condition is naturally satisfied in most economic environments—including ours. (The fact that all branch choice functions in our setting satisfy the irrelevance of rejected contracts condition is immediate from the algorithm defining branch choice—see Lemma D.1.)

²¹As in the work of Hatfield and Milgrom (2005), an irrelevance of rejected contracts condition (which is naturally satisfied in our setting—see Footnote 20) is implicitly assumed throughout the work of Hatfield and Kojima (2010) (Aygün and Sönmez (2012a)).

for all $z, z' \in X$ and $Y \subseteq X$ for which $i(z) \notin i(Y)$ (i.e. no contract in Y is associated to agent $i(z)$).

Unilateral substitutability is a powerful condition, and has been applied in the study of cadet–branch matching mechanisms (Sönmez and Switzer (2011); Sönmez (2011)). Although cadet–branch matching arises as a special case of our framework, the choice functions C^b which arise in markets with slot-specific priorities are not unilaterally substitutable, in general. Our next example illustrates this fact; this also shows (*a fortiori*) that the branch choice functions in our framework may be non-substitutable.

Example 1. Let $X = \{i_0, i_1, j_1\}$, with $B = \{b\}$, $I = \{i, j\}$, $i(i_0) = i = i(i_1)$ and $i(j_1) = j$.²² If b has two slots, $s_b^1 \triangleright^b s_b^2$, with priorities given by

$$\begin{aligned}\Pi^{s_b^1} &: i_0 \succ \emptyset_{s_b^1}, \\ \Pi^{s_b^2} &: i_1 \succ j_1 \succ \emptyset_{s_b^2},\end{aligned}$$

then C^b fails the unilateral substitutes condition: $j_1 \notin C^b(\{i_1, j_1\})$, but $j_1 \in C^b(\{i_0, i_1, j_1\})$.

Nevertheless, the choice functions C^b do behave substitutably whenever each agent offers at most one contract to b .

Definition. A choice function C^b is **substitutable on no-negotiation contract sets** if

$$z \notin C^b(Y \cup \{z\}) \implies z \notin C^b(Y \cup \{z, z'\})$$

for any $z, z' \in X_b$ and $Y \subseteq X_b$ such that

$$|Y \cup \{z, z'\}| = |i(Y \cup \{z, z'\})|. \quad (1)$$

²²Clearly (in order for $b(\cdot)$ to be well-defined), we must have $b(i_0) = b(i_1) = b(j_1) = b$, as $|B| = 1$.

Proposition 1. *Every branch choice function C^b is substitutable on no-negotiation contract sets.*

Also, the choice functions C^b satisfy the weakest form of the substitutability condition introduced by Hatfield and Kojima (2010).

Definition. A choice function C^b is **bilaterally substitutable** if

$$z \notin C^b(Y \cup \{z\}) \implies z \notin C^b(Y \cup \{z, z'\})$$

for all $z, z' \in X$ and $Y \subseteq X$ with $i(z), i(z') \notin i(Y)$.

Proposition 2. *Every choice function C^b is bilaterally substitutable.*

In addition to illustrating some of the structure underlying the choice functions induced by slot-specific priorities, Proposition 1 is also useful in our proofs. Meanwhile, Proposition 2 is not used directly in the sequel—as with Example 1, we present Proposition 2 only to illustrate the relationship between our work and the conditions introduced in the prior literature.²³

3.4.2 The Law of Aggregate Demand

A number of structural results in two-sided matching theory rely on the following monotonicity condition introduced by Hatfield and Milgrom (2005).²⁴

Definition. A choice function C^b **satisfies the Law of Aggregate Demand** if

$$Y' \supseteq Y \implies |C^b(Y')| \geq |C^b(Y)|.$$

²³Combining Proposition 2 with Theorem 1 of Hatfield and Kojima (2010) can be used to prove the existence of stable outcomes in our setting, although (as observed in Footnote 21) this logic implicitly requires an irrelevance of rejected contracts condition which branch choice functions satisfy (Aygün and Sönmez (2012a)). However, as Hatfield and Kojima (2010) pointed out, the bilateral substitutes condition is not sufficient for the other key results necessary for matching market design, such as the existence of strategy-proof matching mechanisms. To obtain these additional results in our framework, we draw upon structures present in our specific model (see Section 4.1); these structures give rise to a self-contained existence proof, which does not make use of the bilateral substitutability condition.

²⁴Alkan (2002) and Alkan and Gale (2003) introduced a related *cardinal monotonicity* condition.

Unfortunately, like the substitutability and unilateral substitutability conditions, the branch choice functions in our framework may fail to satisfy the law of aggregate demand.

Example 2. Let $X = \{i_0, i_1, j_0\}$, with $B = \{b\}$, $I = \{i, j\}$, $i(i_0) = i = i(i_1)$ and $i(j_0) = j$.²⁵ If b has two slots, $s_b^1 \triangleright^b s_b^2$, with priorities given by

$$\Pi^{s_b^1} : i_0 \succ j_0 \succ \emptyset_{s_b^1},$$

$$\Pi^{s_b^2} : i_1 \succ \emptyset_{s_b^2},$$

then C^b does not satisfy the law aggregate demand:

$$|C^b(\{i_1, j_0\})| = |\{i_1, j_0\}| = 2 > 1 = |\{i_0\}| = |C^b(\{i_0, i_1, j_0\})|.$$

4 Basic Theory

4.1 Associated Agent–Slot Matching Market

We now associate our original market to a (one-to-one) **agent–slot matching market**, in which, *slots*, rather than branches, compete for contracts. The contract set X is extended to the set \tilde{X} defined by

$$\tilde{X} \equiv \{\langle x; s \rangle : x \in X \text{ and } s \in S_{b(x)}\}.$$

Slot priorities $\tilde{\Pi}^s$ over contracts in \tilde{X} exactly correspond to the priorities Π^s over contracts in X :

$$\begin{aligned} \langle x; s \rangle \tilde{\Pi}^s \langle x'; s \rangle &\iff x \Pi^s x'; \\ \emptyset_s \tilde{\Pi}^s \langle x; s' \rangle &\iff [\emptyset_s \Pi^s x \text{ or } s' \neq s]. \end{aligned}$$

²⁵Clearly (in order for $b(\cdot)$ to be well-defined), we must have $b(i_0) = b(i_1) = b(j_0) = b$, as $|B| = 1$.

Meanwhile, the preferences \tilde{P}^i of $i \in I$ over contracts in \tilde{X} respect the order P^i , while using orders of precedence to break ties among slots:

$$\begin{aligned}\langle x; s \rangle \tilde{P}^i \langle x'; s' \rangle &\iff x P^i x' \text{ or } [x = x' \text{ and } s \triangleright^{b(x)} s']; \\ \emptyset_i \tilde{P}^i \langle x; s \rangle &\iff [\emptyset_i P^i x \text{ or } i(x) \neq i].\end{aligned}$$

These extended priorities $\tilde{\Pi}^s$ and preferences \tilde{P}^i induce choice functions over \tilde{X} :

$$\begin{aligned}\tilde{C}^s(\tilde{Y}) &\equiv \max_{\tilde{\Pi}^s} \tilde{Y}; \\ \tilde{C}^i(\tilde{Y}) &\equiv \max_{\tilde{P}^i} \tilde{Y}.\end{aligned}$$

To avoid terminology confusion, we call a set $\tilde{Y} \subseteq \tilde{X}$ a **slot-outcome**. It is clear that slot-outcomes $\tilde{Y} \subseteq \tilde{X}$ correspond to outcomes $Y \subseteq X$ according to the natural projection $\varpi : \tilde{X} \rightarrow X$ defined by

$$\varpi(\tilde{Y}) \equiv \{x : \langle x; s \rangle \in \tilde{Y} \text{ for some } s \in S_{b(x)}\}.$$

Our contract set restriction notation extends naturally to slot-outcomes \tilde{Y} :

$$\tilde{Y}_i = \{\langle y; s \rangle \in \tilde{Y} : i(y) = i\}; \quad \tilde{Y}_s = \{\langle y; s' \rangle \in \tilde{Y} : s' = s\}.$$

Definition. A slot-outcome $\tilde{Y} \subseteq \tilde{X}$ is **slot-stable** if it is

1. **individually rational for agents and slots**— $\tilde{C}^i(\tilde{Y}) = \tilde{Y}_i$ for all $i \in I$ and $\tilde{C}^s(\tilde{Y}) = \tilde{Y}_s$ for all $s \in S$ —and
2. **not blocked at any slot**—there does not exist a **slot-block** $\langle z; s \rangle \in \tilde{X}$ such that $\langle z; s \rangle = \tilde{C}^{i(z)}(\tilde{Y} \cup \{\langle z; s \rangle\})$ and $\langle z; s \rangle = \tilde{C}^s(\tilde{Y} \cup \{\langle z; s \rangle\})$.

Lemma 1. *If $\tilde{Y} \subseteq \tilde{X}$ is slot-stable, then $\varpi(\tilde{Y})$ is stable.*

Theorem 3 of Hatfield and Milgrom (2005) implies that one-to-one matching with contracts markets have stable outcomes. Combining this observation with Lemma 1 gives a second proof that the set of stable outcomes is always nonempty in our framework. In the next section, we refine this observation by focusing on the stable outcome associated to the slot-outcome of agent-optimal slot-stable mechanism for the agent–slot market.

4.2 The Cumulative Offer Process

We now introduce the cumulative offer process for matching with contracts (see Hatfield and Kojima (2010); Hatfield and Milgrom (2005); Kelso and Crawford (1982)), which generalizes the *agent-proposing deferred acceptance algorithm* of Gale and Shapley (1962). We provide an intuitive description of this algorithm here; a more technical statement is given in Appendix A.

Definition. In the **cumulative offer process**, agents propose contracts to branches in a sequence of steps $\ell = 1, 2, \dots$:

Step 1. Some agent $i^1 \in I$ proposes his most-preferred contract, $x^1 \in X_{i^1}$. Branch $b(x^1)$ *holds* x^1 if $x^1 \in C^{b(x^1)}(\{x^1\})$, and *rejects* x^1 otherwise. Set $A_{b(x^1)}^2 = \{x^1\}$, and set $A_{b'}^2 = \emptyset$ for each $b' \neq b(x^1)$; these are the sets of contracts *available* to branches at the beginning of Step 2.

Step 2. Some agent $i^2 \in I$ for whom no contract is currently held by any branch proposes his most-preferred contract which has not yet been rejected, $x^2 \in X_{i^2}$. Branch $b(x^2)$ holds the contracts in $C^{b(x^2)}(A_{b(x^2)}^2 \cup \{x^2\})$ and rejects all other contracts in $A_{b(x^2)}^2 \cup \{x^2\}$; branches $b' \neq b(x^2)$ continue to hold all contracts they held at the end of Step 1. Set $A_{b(x^2)}^3 = A_{b(x^2)}^2 \cup \{x^2\}$, and set $A_{b'}^3 = A_{b'}^2$ for each $b' \neq b(x^2)$.

Step ℓ . Some agent $i^\ell \in I$ for whom no contract is currently held by any branch proposes his most-preferred contract which has not yet been rejected, $x^\ell \in$

X_i^ℓ . Branch $b(x^\ell)$ holds the contracts in $C^{b(x^\ell)}(A_{b(x^\ell)}^\ell \cup \{x^\ell\})$ and rejects all other contracts in $A_{b(x^\ell)}^\ell \cup \{x^\ell\}$; branches $b' \neq b(x^\ell)$ continue to hold all contracts they held at the end of Step $\ell - 1$. Set $A_{b(x^\ell)}^{\ell+1} = A_{b(x^\ell)}^\ell \cup \{x^\ell\}$, and set $A_{b'}^{\ell+1} = A_{b'}^\ell$ for each $b' \neq b(x^\ell)$.

If at any time no agent is able to propose a new contract,²⁶ then the algorithm terminates. The **outcome of the cumulative offer process** is the set of contracts held by branches at the end of the last step before termination.

In the cumulative offer process, agents propose contracts sequentially. Branches accumulate offers, choosing at each step (according to C^b) a set of contracts to hold from the set of all previous offers. The process terminates with no agent wishes to propose contracts—that is, when all agents for whom no contracts are held have proposed all contracts they find acceptable.

Note that we do not explicitly specify the order in which agents make proposals. This is because in our setting, the cumulative offer process outcome is in fact *independent* of the order of proposal.²⁷ An analogous order-independence result is known for settings where priorities induce unilaterally substitutable branch choice functions (Hatfield and Kojima (2010)). However, as our discussion in Section 3.4.1 notes, slot-specific priorities may not induce unilaterally substitutable choice functions. Meanwhile, no order-independence result is known for the general class of bilaterally substitutable preferences.²⁸

Our first main result shows that the cumulative offer process outcome has a natural interpretation: it corresponds to the outcome of the agent-optimal slot-stable slot-outcome in the agent–slot matching market.

Theorem 1. *The slot-outcome of the agent-optimal slot-stable mechanism in the agent–slot matching market corresponds (under projection ϖ) to the outcome of the cumulative offer*

²⁶That is, if all agents for whom no contracts are on hold have proposed all contracts they find acceptable.

²⁷We make this statement concrete in Theorem B.1 of Appendix B.

²⁸Although Hatfield and Kojima (2010) state their cumulative offer process algorithm without attention to the proposal order, they only prove that the choice of proposal order has no impact on outcomes in the case of unilaterally substitutable preferences.

process.

The proof of Theorem 1 proceeds in three steps. First, we show that the contracts “held” by each slot improve (with respect to slot priority order) over the course of the cumulative offer process.²⁹ This observation implies that no contract held by a slot $s \in S$ at some step of the cumulative offer process has higher priority than the contract s holds at the end of the process; it follows that the cumulative offer process outcome Y is the ϖ -projection of a slot-stable slot-outcome \tilde{Y} . Then, we demonstrate that agents (weakly) prefer \tilde{Y} to the agent-optimal slot-stable slot outcome \tilde{Z} (which exists by use Theorem 3 of Hatfield and Milgrom (2005)); this implies that $\tilde{Y} = \tilde{Z}$, proving Theorem 1 as $\varpi(\tilde{Y}) = Y$.

Theorem 1 implies that the cumulative offer process always terminates.³⁰ Moreover, it shows (by Lemma 1) that the cumulative offer process outcome is stable and somewhat distinguished among stable outcomes.

Theorem 2. *The cumulative offer process produces an outcome which is stable. Moreover, for any slot-stable $\tilde{Z} \subseteq \tilde{X}$, each agent (weakly) prefers the outcome of the cumulative offer process to $\varpi(\tilde{Z})$.*

Note that Theorem 2 shows only that agents weakly prefer the cumulative offer process outcome to any other stable outcome associated to a slot-stable slot-outcome. As not all stable outcomes are associated to slot-stable slot-outcomes, this need not imply that each agent prefers the cumulative offer process outcome to all other stable outcomes; we demonstrate this explicitly in the next section.

4.3 Agent-Optimal Stable Outcomes

We say that an outcome $Y \subseteq X$ **Pareto dominates** $Y' \subseteq X$ if $Y_i R^i Y'_i$ for all $i \in I$, and $Y_i P^i Y'_i$ for at least one $i \in I$. An outcome $Y \subseteq X$ which Pareto dominates all other stable

²⁹Here, by the contract “held” by a slot $s \in S_b$ in step ℓ , we mean the contract assigned to s in the computation of $C^b(A_b^{\ell+1})$.

³⁰This fact can also be observed directly, as the set X is finite, and the available contract sets A_b^ℓ grow monotonically in ℓ .

outcomes is called an **agent-optimal stable outcome**.³¹ For general slot-specific priorities, agent-optimal stable outcomes need not exist, as the following example shows.

Example 3. Let $X = \{i_0, i_1, j_0, j_1, k_0, k_1\}$, with $B = \{b\}$, $I = \{i, j, k\}$ and $i(h_0) = h = i(h_1)$ for each $h \in I$.³² We suppose that $h_0 P^h h_1 P^h \emptyset_h$ for each $h \in I$, and that b has two slots, $s_b^1 \triangleright^b s_b^2$, with slot priorities given by

$$\begin{aligned}\Pi^{s_b^1} &: i_1 \succ j_1 \succ k_1 \succ i_0 \succ j_0 \succ k_0 \succ \emptyset_{s_b^1}, \\ \Pi^{s_b^2} &: i_0 \succ i_1 \succ j_0 \succ j_1 \succ k_0 \succ k_1 \succ \emptyset_{s_b^2}.\end{aligned}$$

In this setting, the outcomes $Y \equiv \{j_1, i_0\}$ and $Y' \equiv \{i_1, j_0\}$ are both stable. However, $Y_i P^i Y'_i$ while $Y_j P^j Y'_j$, so there is no agent-optimal stable outcome.

Here, Y is associated to a slot-stable slot-outcome, but Y' is not. As we expect from Theorem 2, the cumulative offer process produces the former of these two outcomes, Y .

Although matching markets with slot-specific priorities may not have agent-optimal stable outcomes, the cumulative offer process finds agent-optimal stable outcomes when they exist.

Theorem 3. *If an agent-optimal stable outcome exists, then it is the outcome of the cumulative offer process.*

In our proof of Theorem 3, we show that no stable outcome can Pareto dominate the cumulative offer process outcome. That is: for any stable outcome Y which is not equal to the outcome Z of the cumulative offer process, there is some agent $i \in I$ such that $Z_i P^i Y_i$. This quickly implies Theorem 3, as agent-optimal stable outcomes Pareto dominate all other stable outcomes.

³¹That is, an agent-optimal stable outcome is a stable outcome such that $Y_i R^i Y'_i$ for any agent $i \in I$ and stable outcome $Y' \subseteq X$.

³²Clearly (in order for $b(\cdot)$ to be well-defined), we must have $b(h_0) = b = b(h_1)$ for each $h \in I$, as $|B| = 1$.

4.4 The Cumulative Offer Mechanism

A **mechanism** consists of a **strategy space** \mathcal{S}^i for each agent $i \in I$, along with an **outcome function** $\varphi_{\Pi} : \prod_{i \in I} \mathcal{S}^i \rightarrow X$ that selects an outcome for each choice of agent strategies. We confine our attention to **direct mechanisms**, i.e. mechanisms for which the strategy spaces correspond to the preference domains: $\mathcal{S}^i \equiv \mathcal{P}^i$, where \mathcal{P}^i denotes the set of all possible preference relations for agent $i \in I$. Such mechanisms are entirely determined by their outcome functions, hence in the sequel we identify mechanisms with their outcome functions and use the term “mechanism φ_{Π} ” to refer to the mechanism with outcome function φ_{Π} and $\mathcal{S}^i \equiv \mathcal{P}^i$ (for all $i \in I$). All mechanisms we discuss implicitly depend on the priority profile under consideration; we often suppress the priority profile from the mechanism notation, writing “ φ ” instead of “ φ_{Π} ,” if doing so will not introduce confusion.

In this section, we undertake analysis of the **cumulative offer mechanism (associated to slot priorities Π)**, which selects the outcome obtained by running the cumulative offer process (with respect to submitted preferences and priorities Π). We denote this mechanism by $\Phi_{\Pi} : \prod_{i \in I} \mathcal{P}^i \rightarrow X$.

4.4.1 Stability and Strategy-Proofness of Φ_{Π}

A mechanism φ is **stable** if it always selects an outcome stable with respect to the input preferences.

We say that a mechanism φ is **strategy-proof** if there is no agent $i \in I$, $P^i \in \prod_{i \in I} \mathcal{P}^i$, and $\bar{P}^i \neq P^i$ such that $\varphi(\bar{P}^i, P^{-i}) P^i \varphi(P^i, P^{-i})$. Similarly, we say that a mechanism φ is **group strategy-proof** if there is no set of agents $I' \subset I$, $P^{I'} \in \prod_{i \in I'} \mathcal{P}^i$, and $\bar{P}^{I'} \neq P^{I'}$ such that $\varphi(\bar{P}^{I'}, P^{-I'}) P^i \varphi(P^{I'}, P^{-I'})$ for all $i \in I'$.

It follows immediately from Theorem 2 that the cumulative offer mechanism is stable. Meanwhile, Theorem 1 of Hatfield and Kojima (2009) implies that the agent-optimal slot-stable mechanism is (group) strategy-proof in the agent–slot matching market. Thus, we see that the cumulative offer mechanism is (group) strategy-proof, as any $\bar{P}^{I'} \neq P^{I'}$ such

that $\varphi(\bar{P}^{I'}, P^{-I'}) P^i \varphi(P^{I'}, P^{-I'})$ for all $i \in I'$ would give rise to a profitable manipulation ($\tilde{P}^{I'} \neq \bar{P}^{I'}$) of the agent-optimal slot-stable mechanism. These observations are summarized in the following theorem.

Theorem 4. *The cumulative offer mechanism Φ_{Π} is*

1. *stable, and*
2. *(group) strategy-proof.*

4.4.2 Φ_{Π} Respects Unambiguous Improvements

We say that priority profile $\bar{\Pi}$ is an **unambiguous improvement over priority profile Π for $i \in I$** if

1. for all $x \in X_i$ and $y \in X_{I \setminus \{i\}}$, if $x \Pi^s y$, then $x \bar{\Pi}^s y$; and
2. for all $y, z \in X_{I \setminus \{i\}}$, $y \Pi^s z$ if and only if $y \bar{\Pi}^s z$.

That is, $\bar{\Pi}$ is an unambiguous improvement over priority profile Π for $i \in I$ if $\bar{\Pi}$ is obtained from Π by increasing the priorities of some of i 's contracts (at some slots) while leaving the relative priority orders of other agents' contracts unchanged.

We say that a mechanism φ **respects unambiguous improvements for i** if for any preference profile P^I ,

$$(\varphi_{\bar{\Pi}}(P^I))_i R^i (\varphi_{\Pi}(P^I))_i$$

whenever $\bar{\Pi}$ is an unambiguous improvement over Π for i . And we say that φ **respects unambiguous improvements** if it respects unambiguous improvements for each agent $i \in I$.

Theorem 5. *The cumulative offer mechanism Φ_{Π} respects unambiguous improvements.*

Our proof of Theorem 5 makes use of the fact that the cumulative offer process outcome is independent of the contract proposal order. In particular, we focus on a proposal order

in which i proposes contracts only when no other agent is able to propose. This choice of proposal order guarantees that i is always the last agent to propose a contract in the running of the cumulative offer process (for any priority profile). As $\bar{\Pi}$ is an unambiguous improvement over Π for i , we can show that the last contract i proposes in the cumulative offer process with priority profile $\bar{\Pi}$ must also be proposed in the cumulative offer process with priority profile Π . This yields the desired result because it implies that i is at least as well off under the outcome of cumulative offer process with priority profile $\bar{\Pi}$ as under the cumulative offer process with priority profile Π .

As an alternative to this approach to the proof of Theorem 5, we could instead show that the agent-optimal slot-stable mechanism for the agent–slot matching market satisfies a condition analogous to respecting unambiguous improvements. Theorem 5 would then follow from Theorem 1.

One note: A natural strengthening of our notion of an unambiguous improvement for $i \in I$ would include the condition that i 's preferred contracts (weakly) increase in priority—formally, for all $b \in B$, $s \in S_b$, and $x, x' \in (X_i \cap X_b)$,

$$\text{if } x\bar{\Pi}^s x' \text{ and } x'\Pi^s x, \text{ then } xP^i x'. \tag{2}$$

As Theorem 5 shows that Φ_{Π} respects unambiguous improvements, we see *a fortiori* that Φ_{Π} respects unambiguous improvements that satisfy the additional condition (2).

5 Further Applications

5.1 Design of Affirmative Action Mechanisms

We say that a matching problem has **agent types** if the contract set X is a subset of $I \times B \times \mathbb{T}$ for some **type set** \mathbb{T} , and for each $i \in I$, $X_i = \{i\} \times B \times \{t\}$ for some $t \in \mathbb{T}$,

so that each i is associated to exactly one **type** \mathbf{t} .³³ For such a problem, we identify agents with their types, writing $\mathbf{t}(i)$ for the unique type $\mathbf{t} \in \mathbf{T}$ such that $X_i = \{i\} \times B \times \{\mathbf{t}\}$. For consistency with the prior literature on school choice, we abuse notation slightly by writing i to denote, for each branch $b \in B$, the unique contract $(i, b, \mathbf{t}(i)) \in (X_i \cap X_b)$.

Imposing this additional structure on our general model simplifies the form of slot-specific priorities, rendering branches' choice functions substitutable.

Proposition 3. *In a matching problem with slot-specific priorities and agent types, the branch choice functions C^b are substitutable and satisfy the law of aggregate demand.*

In settings with agent types, substitutability coincides with substitutability on no-negotiability contract sets—this conclusion obtains whenever $|X_i \cap X_b| \leq 1$ for all $i \in I$ and $b \in B$. Thus, Proposition 3 follows directly from Proposition 1.

Combining Proposition 3 with Theorems 15, 3, and 4 of Hatfield and Milgrom (2005) shows that there exists an agent-optimal stable outcome in any matching problem with slot-specific priorities and agent types. The following result then follows upon combining this observation with our Theorems 3, 4, and 5.

Corollary 1. *In a matching problem with slot-specific priorities and agent types, the cumulative offer mechanism Φ_Π is an agent-optimal stable mechanism, which is (group) strategy-proof and respects unambiguous improvements.*

As our discussion in Section 2.1 suggests, decreases in the precedence of slots that rank agents of type \mathbf{t} highly can improve type- \mathbf{t} agents' cumulative offer outcomes. Unfortunately, while we believe this comparative static should hold in reasonably-sized marketplaces, it may fail in small markets.³⁴

³³Note that we may assume without loss of generality that $X_i = \{i\} \times B \times \{\mathbf{t}\}$, as any case in which $X_i \subsetneq \{i\} \times B \times \{\mathbf{t}\}$ can then be captured by assuming some contracts $x \in \{i\} \times B \times \{\mathbf{t}\}$ to be unacceptable to slots at their associated branches $\mathbf{b}(x)$.

³⁴Example C.1 of Appendix C illustrates this fact.

5.1.1 “Soft” Minority Quotas in School Choice

Many affirmative action programs impose quotas on majority agents. However, as Kojima (2012) showed, quota policies can have perverse effects: some quota-based affirmative action policies hurt *all* minority students under any stable matching mechanism. Despite these discouraging observations, Hafalir et al. (forthcoming) recently introduced a novel approach to affirmative action, *affirmative action with minority reserves*, which compares favorably to the more standard majority-quota policies.

In the Hafalir et al. (forthcoming) approach, certain slots at each school are *reserved* for minorities but convert into regular slots if not claimed by minority students. Formally, the model of Hafalir et al. (forthcoming) embeds into the framework of matching with slot-specific priorities as follows: The agents $i \in I$ are *students* and the branches $b \in B$ are *schools*. Each student $i \in I$ has a strict linear preference order P^i over schools, and is of either minority (m) or majority (M) type (i.e. $t(i) \in \{\mathbf{M}, \mathbf{m}\} = \mathbf{T}$). Each school $b \in B$ has a strict linear order π^b over students and a number of slots q_b corresponding to its “capacity.”

Under **affirmative action with minority reserves**, each school $b \in B$ has an associated minority reserve $r_b^m \leq q_b$ such that b prefers any minority applicant to any majority applicant if the number of minority students admitted to is below r_b^m (Hafalir et al. (forthcoming)). This policy can be implemented by choosing slot-specific priorities $\bar{\Pi}$ such that

1. for all $\ell \leq r_b^m$, $i\bar{\Pi}^{s_b^\ell}i'\bar{\Pi}^{s_b^\ell}\emptyset_{s_b^\ell} \iff$
 - (a) $t(i) = \mathbf{m}$ and $t(i') = \mathbf{M}$, or
 - (b) $t(i) = t(i')$ and $i\pi^b i'$;
2. for all $\ell > r_b^m$, $i\bar{\Pi}^{s_b^\ell}i'\bar{\Pi}^{s_b^\ell}\emptyset_{s_b^\ell} \iff i\pi^b i'$.

Under **affirmative action with majority quotas**, meanwhile, each school $b \in B$ has an associated majority quota $q_b^M \leq q_b$ such that b cannot admit more than q_b^M majority applicants. This policy can be implemented by choosing slot-specific priorities Π such that

1. for all $\ell < q_b - q_b^M$, $i\Pi^{s_b^\ell}i'\Pi^{s_b^\ell}\emptyset_{s_b^\ell} \iff t(i) = t(i') = m$ and $i\pi^b i'$;
2. for all $\ell \geq q_b - q_b^M$, $i\Pi^{s_b^\ell}i'\Pi^{s_b^\ell}\emptyset_{s_b^\ell} \iff i\pi^b i'$.

With these observations, we may derive two of the main results of Hafalir et al. (forthcoming) as consequences of our general results for slot-specific priority structures.

Proposition 4 (Hafalir et al. (forthcoming)). *1. In the presence of affirmative action with minority reserves, the cumulative offer mechanism produces the student-optimal stable outcome and is (group) strategy-proof.*

2. Given a vector q^M of majority quotas, set $r_b^m = q_b - q_b^M$ for each $b \in B$, and let Y be an outcome which is stable under the priorities Π induced by the quotas q^M and tiebreakers π . Either:

- (a) Y is stable under the priorities $\bar{\Pi}$ induced by reserves r^m and tiebreakers π , or*
- (b) there exists an outcome Z which is stable under priorities $\bar{\Pi}$ and Pareto dominates Y .*

Proof. Part 1 follows directly from Proposition 3. Part 2 also follows quickly: By Corollary 1, we know that $(\Phi_\Pi(P^I))_i R^i Y_i$ for each $i \in I$. Meanwhile, $\bar{\Pi}$ is an unambiguous (weak) improvement for each $i \in I$, hence $(\Phi_{\bar{\Pi}}(P^I))_i R^i (\Phi_\Pi(P^I))_i$ for each $i \in I$ (by Theorem 5). Thus, taking $Z \equiv \Phi_{\bar{\Pi}}(P^I)$ gives

$$Z_i = (\Phi_{\bar{\Pi}}(P^I))_i R^i (\Phi_\Pi(P^I))_i R^i Y_i, \quad (3)$$

for each $i \in I$. Now, if $Y = \Phi_\Pi(P^I) = \Phi_{\bar{\Pi}}(P^I) = Z$, then Y is stable under the priorities $\bar{\Pi}$. Otherwise, there is at least one $i \in I$ for whom the identity (3) is strict. In that case, Z is stable under priorities $\bar{\Pi}$ and Pareto dominates Y . \square

$$\left\{ \begin{array}{ll} 4 \sim 3 \sim 2 \sim 1 \succ \emptyset & 0 < \ell \leq \frac{40}{100}q_b \\ 4 \succ \emptyset & \frac{40}{100}q_b < \ell \leq \frac{55}{100}q_b \\ 3 \succ \emptyset & \frac{55}{100}q_b < \ell \leq \frac{70}{100}q_b \\ 2 \succ \emptyset & \frac{70}{100}q_b < \ell \leq \frac{85}{100}q_b \\ 1 \succ \emptyset & \frac{85}{100}q_b < \ell \leq \frac{100}{100}q_b. \end{array} \right.$$

Figure 2: Structure of slot priorities in the Chicago selective high school match. The top 40% of slots are *open slots*; the bottom 60% feature *socioeconomic reserves* for each of the four student types ($\{4, 3, 2, 1\}$).

5.1.2 Socioeconomic Affirmative Action in Chicago School Choice

We now demonstrate that our framework embeds the real-world structure of the Chicago selective high school choice affirmative action program discussed in Section 2.1. Here, the agents $i \in I$ and branches $b \in B$ again correspond to *students* and *schools*. Each student $i \in I$ has a strict linear preference order P^i over schools, and there are four agent types representing the different SES tiers: $\mathbb{T} = \{4, 3, 2, 1\}$.

The top 40% of the q_b slots at school $b \in B$ are **open slots**, assigned based on a strict linear merit order π^* over students, which is determined by composite test scores—and thus uniform across schools. The remaining 60% of the slots at each school $b \in B$ feature **socioeconomic reserves**: the first 15% of these slots are reserved for students $i \in I$ of type $t(i) = 4$; the next 15% are reserved for students $i \in I$ of type $t(i) = 3$; and so forth. These type priority structures are illustrated in Figure 2.

Formally, the set S_b of slots at school b is partitioned into subsets

$$S_b = S_b^\circ \cup S_b^4 \cup S_b^3 \cup S_b^2 \cup S_b^1,$$

with S_b° consisting of $\frac{40}{100}q_b$ slots, and each set S_b^t consisting of $\frac{15}{100}q_b$ slots.³⁵ The priorities of slots $s \in S_b^\circ$ are such that

$$i \Pi^s i' \Pi^s \emptyset_s \iff i \pi^* i'.$$

³⁵We assume for simplicity that q_b is a multiple of 20, so that $\frac{15}{100}q_b, \frac{40}{100}q_b \in \mathbb{Z}$.

$$\left\{ \begin{array}{ll} 4 \succ 3 \sim 2 \sim 1 \succ \emptyset & 0 < \ell \leq \frac{15}{100}q_b \\ 3 \succ \emptyset & \frac{15}{100}q_b < \ell \leq \frac{30}{100}q_b \\ 2 \succ \emptyset & \frac{30}{100}q_b < \ell \leq \frac{45}{100}q_b \\ 1 \succ \emptyset & \frac{45}{100}q_b < \ell \leq \frac{60}{100}q_b \\ 4 \sim \emptyset & \frac{60}{100}q_b < \ell \leq \frac{100}{100}q_b. \end{array} \right.$$

Figure 3: Counterfactual slot priority structure for the Chicago selective high school match. Here, the top 60% of slots feature socioeconomic reserves, while the bottom 40% are open.

Meanwhile, the priorities of slots $s \in S_b^t$ are such that

$$i\Pi^s i' \Pi^s \emptyset_s \iff i\pi^* i'$$

whenever $t(i) = t(i') = \mathbf{t}$, and $\emptyset_s \Pi^s i$ whenever $t(i) \neq \mathbf{t}$.³⁶ The order of precedence \triangleright^b is such that

$$s^\circ \triangleright^b s^4 \triangleright^b s^3 \triangleright^b s^2 \triangleright^b s^1 \tag{4}$$

for all $s^\circ \in S_b^\circ$, $s^4 \in S_b^4$, $s^3 \in S_b^3$, $s^2 \in S_b^2$, and $s^1 \in S_b^1$.³⁷

Tables 1–3 show the effect of switching from the current precedence orders \triangleright^b to the alternate orders, illustrated in Figure 3, in which open slots are filled after reserve slots. Formally, these counterfactual precedence orders $\bar{\triangleright}^b$ are such that

$$s^4 \bar{\triangleright}^b s^3 \bar{\triangleright}^b s^2 \bar{\triangleright}^b s^1 \bar{\triangleright}^b s^\circ$$

for all $s^\circ \in S_b^\circ$, $s^4 \in S_b^4$, $s^3 \in S_b^3$, $s^2 \in S_b^2$, and $s^1 \in S_b^1$.

Our model suggests a natural precedence order which gives rise to priorities in between the current CPS priority structure—under which all open slots are filled first—and the counterfactual structure discussed in Section 2.1—under which all reserved slots are filled

³⁶Because all the slots in Chicago’s selective high schools are overdemand, all seats reserved for students of type \mathbf{t} are claimed by students in type \mathbf{t} , hence we may assume for expositional simplicity that slots $s \in S_b^t$ find students of types $\mathbf{t}' \neq \mathbf{t}$ unacceptable.

³⁷Since all slots $s^e \in S_b^e$ have identical priorities, (4) suffices to specify the precedence order up to equivalence.

first. Instead of filling all the open slots at once, CPS could alternate between filling open slots and filling reserve slots, in proportion to the total numbers of each slot type available. For example, **intermediate priorities** could be designed so as to fill three open slots at each school, then one of each type of reserved slot at each school, then three more open slots, then four more reserved slots, and so forth.³⁸ This approach spreads the access to open seats evenly throughout the priority structure, virtually eliminating the biases of the current CPS system.³⁹

Simulation results presented in Table 4 shows that student outcomes under the intermediate priorities are almost identical to those arising when all reserved seats are filled before the open seats (the counterfactual discussed in Section 2.1). While this might at first seem surprising, is in fact quite natural, given the distribution of CPS students' test scores. As we pointed out in Section 2.1, high-SES students' composite scores dominate low-SES students' scores throughout the relevant part of the score distribution. As a result, high-SES students fill the first available open seats, then the highest-scoring low-SES students receive the first reserved seats. Then, once again, the top of the truncated scoring distribution consists of high-SES students; the students take the next open slots. Once again, then, the highest-scoring low-SES students who remain unassigned receive reserved seats, leaving even fewer low-SES students with high scores in the pool. This process produces an outcome very similar to that found when all reserved seats are filled before the open seats.

5.2 Precedence Order Changes in Cadet–Branch Matching

The cadet–branch matching problem studied by Sönmez and Switzer (2011) and Sönmez (2011) is a matching problem with $X = I \times B \times \{t_0, t_+\}$. Here, the agents $i \in I$ correspond to *cadets*, who must be assigned to *branches of service* $b \in B$. Contracts with term t_0 represent standard service contracts; contracts with term t_+ represent the standard contract

³⁸When using this approach, every third block of slots should have only two open slots, so as to maintain the overall 40%-15%-15%-15%-15% proportions in every block of 20 slots.

³⁹As the number of slots at each school is finite, completely eliminating the bias would require randomizing the precedence relation to some extent.

	Intermediate Mechanism (Open Slots First in Group)				Effect of Switching (to fill Open Slots Last)			
	Tier 4	Tier 3	Tier 2	Tier 1	Tier 4	Tier 3	Tier 2	Tier 1
1	135	53	39	39	4	-2	-1	-1
2	95	138	56	50	2	-2	0	0
3	51	94	121	58	-5	4	3	-2
4	85	102	108	52	-2	1	1	0
5	230	98	69	69	-1	2	-1	0
6	146	54	42	42	1	-1	0	0
7	684	449	253	244	-4	3	1	0
8	93	54	31	31	0	0	0	0
9	65	134	109	97	-3	0	0	3
TOT	1584	1176	828	682	-8	5	3	0

Table 4: Comparison between the unbiased, intermediate mechanism for CPS selective high school enrollment and the counterfactual discussed in Section 2.1, in which open slots are filled last.

supplemented with a three-year service extension.

Cadets are ranked according to a strict linear *order of merit* ranking π^* . The slots of each branch $b \in B$ are partitioned into two sets, $S_b^0 \subset S_b$ and $S_b^+ \subset S_b$, of sizes $(1 - \lambda)q_b$ and λq_b , respectively. Slots $s \in S_b^0$ are **regular slots**, whose priority rankings follow the order of merit list exactly: for any $b \in B$, $s \in S_b^0$, $i \neq i' \in I$, and $\mathbf{t}, \mathbf{t}' \in \{\mathbf{t}_0, \mathbf{t}_+\}$,

$$(i, b, \mathbf{t})\Pi^s(i', b, \mathbf{t}') \iff i\pi^*i'.$$

Our results are independent of how regular slots' relative priorities of contracts (i, b, \mathbf{t}) and (i, b, \mathbf{t}') are specified; for concreteness, we follow the military's convention of assuming that

$$(i, b, \mathbf{t}_0)\Pi^s(i, b, \mathbf{t}_+)$$

for all slots $s \in S_b^0$. Slots $s \in S_b^+$ are **branch-of-choice slots**, which give priority to \mathbf{t}_+ contracts: for any $b \in B$, $s \in S_b^+$, and $i \neq i' \in I$,

$$(i, b, \mathbf{t}_+)\Pi^s(i', b, \mathbf{t}_0), \quad \text{and}$$

$$(i, b, \mathbf{t})\Pi^s(i', b, \mathbf{t}) \iff i\pi^*i'$$

for any $\mathbf{t} \in \{\mathbf{t}_0, \mathbf{t}_+\}$.

In the settings of Sönmez and Switzer (2011) and Sönmez (2011), the branch-of-choice slots have lowest precedence at each branch. That is, the slots $s, s' \in S_b$ at branch $b \in B$ follow a precedence order \blacktriangleright^b such that

$$s \in S_b^0 \text{ and } s' \in S_b^+ \implies s \blacktriangleright^b s'. \quad (5)$$

In our model, all precedence orders satisfying condition (5) are equivalent; hence, we identify the full class of such orders with a “single” precedence order \blacktriangleright^b .

Sönmez and Switzer (2011) demonstrated that the branch choice functions induced by precedence order \blacktriangleright^b are unilaterally substitutable. This implies the existence of a cadet-optimal stable outcome, and allowed Sönmez and Switzer (2011) to propose the use of a cadet-optimal stable mechanism for cadet–branch matching.⁴⁰

Our next result shows that the structure found by Sönmez and Switzer (2011) is unique to the specific precedence order the United States military selected: up to equivalence, \blacktriangleright^b is the only precedence order which guarantees the existence of cadet-optimal stable outcomes in general.

Proposition 5. *For any cadet–branch matching problem precedence order $\triangleright \neq \blacktriangleright$ for which there exists $b \in B$, $s \in S_b^0$, and $s' \in S_b^+$ such that $s' \triangleright^b s$, there exist preference orders for all cadets such that no outcome stable with respect to the branch choice functions C^b induced by the slot priorities Π^s ($s \in S_b$) and precedence order \triangleright^b is cadet-optimal.*

References

Abdulkadiroğlu, Atila and Tayfun Sönmez, “School Choice: A Mechanism Design Approach,” *American Economic Review*, 2003, 93, 729–747.

⁴⁰In his discussion of market design for the ROTC cadet–branch match, Sönmez (2011) extended these results to the case in which more than two distinct contract terms are available.

- , **Parag A. Pathak, Alvin E. Roth, and Tayfun Sönmez**, “The Boston public school match,” *American Economic Review*, 2005, *95*, 368–371.
- , – , – , **and** – , “Changing the Boston School Choice Mechanism: Strategy-proofness as Equal Access,” 2006. Mimeo, Massachusetts Institute of Technology.
- , – , **and** – , “The New York City high school match,” *American Economic Review*, 2005, *95*, 364–367.
- , – , **and** – , “Strategyproofness versus Efficiency in Matching with Indifferences: Redesigning the NYC High School Match,” *American Economic Review*, 2009, *99* (5), 1954–1978.
- Adachi, H.**, “On a characterization of stable matchings,” *Economics Letters*, 2000, *68*, 43–49.
- Alkan, Ahmet**, “A class of multipartner matching markets with a strong lattice structure,” *Economic Theory*, 2002, *19*, 737–746.
- **and David Gale**, “Stable Schedule Matching under Revealed Preference,” *Journal of Economic Theory*, 2003, *112*, 289–306.
- Aygün, Orhan and Tayfun Sönmez**, “The Importance of Irrelevance of Rejected Contracts in Matching under Weakened Substitutes Conditions,” 2012. Mimeo, Boston College.
- **and** – , “Matching with Contracts: The Critical Role of Irrelevance of Rejected Contracts,” 2012. Mimeo, Boston College.
- Balinski, Michel and Tayfun Sönmez**, “A tale of two mechanisms: Student placement,” *Journal of Economic Theory*, 1999, *84*, 73–94.
- Dubins, L. E. and D. A. Freedman**, “Machiavelli and the Gale-Shapley Algorithm,” *American Mathematical Monthly*, 1981, *88*, 485–494.

- Echenique, Federico**, “Contracts vs. Salaries in Matching,” *American Economic Review*, 2012, *102*, 594–601.
- **and Jorge Oviedo**, “Core many-to-one matchings by fixed-point methods,” *Journal of Economic Theory*, 2004, *115*, 358–376.
- Fleiner, T.**, “A fixed point approach to stable matchings and some applications,” *Mathematics of Operations Research*, 2003, *28*, 103–126.
- Friedman, Milton**, “The Role of Government in Education,” in Robert A. Solo, ed., *Economics and the Public Interest*, Rutgers University Press 1955, pp. 123–144.
- , *Capitalism and Freedom*, University of Chicago Press, 1962.
- Gale, David and Lloyd S. Shapley**, “College admissions and the stability of marriage,” *American Mathematical Monthly*, 1962, *69*, 9–15.
- Hafalir, Isa Emin, M. Bumin Yenmez, and Muhammed Ali Yildirim**, “Effective Affirmative Action in School Choice,” *Theoretical Economics*, forthcoming.
- Hatfield, John William and Fuhito Kojima**, “Matching with Contracts: Comment,” *American Economic Review*, 2008, *98*, 1189–1194.
- **and** – , “Group Incentive Compatibility for Matching with Contracts,” *Games and Economic Behavior*, 2009, *67*, 745–749.
- **and** – , “Substitutes and Stability for Matching with Contracts,” *Journal of Economic Theory*, 2010, *145*, 1704–1723.
- **and Paul Milgrom**, “Matching with Contracts,” *American Economic Review*, 2005, *95*, 913–935.
- **and Scott Duke Kominers**, “Contract Design and Stability in Matching Markets,” 2010. Mimeo, Harvard Business School.

- **and** –, “Matching in Networks with Bilateral Contracts,” *American Economic Journal: Microeconomics*, 2012, *4*, 176–208.
- , **Fuhito Kojima, and Yusuke Narita**, “Promoting School Competition Through School Choice: A Market Design Approach,” 2012. Mimeo, Stanford Graduate School of Business.
- Kelso, Alexander S. and Vincent P. Crawford**, “Job Matching, Coalition Formation, and Gross Substitutes,” *Econometrica*, 1982, *50*, 1483–1504.
- Kojima, Fuhito**, “School Choice: Impossibilities for Affirmative Action,” *Games and Economic Behavior*, 2012, *75*, 685–693.
- Kominers, Scott Duke**, “On the Correspondence of Contracts to Salaries in (Many-to-Many) Matching,” *Games and Economic Behavior*, 2012, *75*, 984–989.
- Ostrovsky, Michael**, “Stability in supply chain networks,” *American Economic Review*, 2008, *98*, 897–923.
- Pathak, Parag A. and Tayfun Sönmez**, “Leveling the Playing Field: Sincere and Sophisticated Players in the Boston Mechanism,” *The American Economic Review*, 2008, *98*, 1636–1652.
- **and** –, “School Admissions Reform in Chicago and England: Comparing Mechanisms by Their Vulnerability to Manipulation,” *American Economic Review*, forthcoming.
- Roth, Alvin E.**, “The economics of matching: Stability and incentives,” *Mathematics of Operations Research*, 1982, *7*, 617–628.
- , “The Evolution of the Labor Market for Medical Interns and Residents: A Case Study in Game Theory,” *Journal of Political Economy*, 1984, *92*, 991–1016.
- **and Elliott Peranson**, “The effects of the change in the NRMP matching algorithm,” *American Economic Review*, 1999, *89*, 748–780.

Sönmez, Tayfun, “Bidding for Army Career Specialties: Improving the ROTC Branching Mechanism,” 2011. Mimeo, Boston College.

– **and Tobias B. Switzer**, “Matching with (Branch-of-Choice) Contracts at United States Military Academy,” 2011. Mimeo, Boston College.

A Formal Description of the Cumulative Offer Process

The **cumulative offer process associated to proposal order** \sqsupset is the following algorithm:

1. Let $\ell = 0$. For each $b \in B$, let $D_b^0 \equiv \emptyset$, and let $A_b^1 \equiv \emptyset$.
2. For each $\ell = 1, \dots, |X|$:
 - (a) Let i be the \sqsupset_ℓ -maximal agent $i \in I$ such that $i \notin i(\cup_{b \in B} D_b^{\ell-1})$ and $\max_{P^i}(X \setminus (\cup_{b \in B} A_b^\ell))_i \neq \emptyset_i$ —that is, the agent highest in the proposal order who wants to propose a new contract—if such an agent exists. (If no such agent exists, then proceed to Step 3, below.)
 - i. Let $x \equiv \max_{P^i}(X \setminus (\cup_{b \in B} A_b^\ell))_i$ be i 's most-preferred contract that has not yet been proposed.
 - ii. Let $b \equiv b(x)$. Set $D_b^\ell = C^b(A_b^\ell \cup \{x\})$ and set $A_b^{\ell+1} = A_b^\ell \cup \{x\}$. For each $b' \neq b$, set $D_{b'}^\ell = D_{b'}^{\ell-1}$ and set $A_{b'}^{\ell+1} = A_{b'}^\ell$.

3. Return the outcome

$$Y \equiv \left(\bigcup_{b \in B} D_b^{\ell-1} \right) = \left(\bigcup_{b \in B} C^b(A_b^\ell) \right)$$

consisting of contracts held by branches at the point when no agent wants to propose additional contracts.

Here, the sets $D_b^{\ell-1}$ and A_b^ℓ denote the sets of contracts **held by** and **available to** branch b at the beginning of cumulative offer process step ℓ . We say that a contract z is **rejected** during the cumulative offer process if $z \in A_{b(z)}^\ell$ but $z \notin D_{b(z)}^{\ell-1}$ for some ℓ .

B Proofs Omitted from the Main Text

Proof of Proposition 1

We prove the following auxiliary lemma which directly implies Proposition 1.

Lemma B.1. *Suppose that $Y \subseteq Y' \subseteq X_b$, $|Y| = |i(Y)|$, and $|Y'| = |i(Y')|$. Then, if $y \in Y$ and $y' \in Y'$ are the contracts assigned to $s \in S_b$ in the computations of $C^b(Y)$ and $C^b(Y')$, respectively, we have $y' \Gamma^s y$.*

Proof. The hypotheses on Y and Y' imply that $Y_{i(x)} = \{x\}$ for each $x \in Y$ and that $Y'_{i(x')} = \{x'\}$ for each $x' \in Y'$. With this observation, the following claim follows quickly.

Claim. *Let $V_b^\ell(Z)$ denote the set V_b^ℓ defined in step $\ell-1$ of the computation of $C^b(Z)$. Then, $V_b^\ell(Y) \subseteq V_b^\ell(Y')$.*

Proof. We proceed by induction. We have $V_b^1(Y) = V_b^1(Y')$ *a priori*, so we assume that $V_b^{\ell'}(Y) \subseteq V_b^{\ell'}(Y')$ for all $\ell' < \ell + 1$ for some $\ell > 0$. We now show that this hypothesis implies that $V_b^{\ell+1}(Y) \subseteq V_b^{\ell+1}(Y')$: Let $x' \equiv \max_{\Pi^{s_b^\ell}}(V_b^\ell(Y'))$. If $x' \in V_b^{\ell+1}(Y)$, then clearly $x' = \max_{\Pi^{s_b^\ell}}(V_b^\ell(Y))$; hence,

$$V_b^{\ell+1}(Y) = (V_b^\ell(Y)) \setminus Y_{i(x')} = (V_b^\ell(Y)) \setminus \{x'\} \subseteq (V_b^\ell(Y')) \setminus \{x'\} = (V_b^\ell(Y')) \setminus Y'_{i(x')} = V_b^{\ell+1}(Y')$$

as desired. Otherwise, we have $x' \notin V_b^{\ell+1}(Y)$, so that $\left(\max_{\Pi^{s_b^\ell}}(V_b^\ell(Y))\right) \equiv x \neq x'$. As $x \in V_b^\ell(Y) \subseteq V_b^\ell(Y') \setminus \{x'\}$, we have

$$V_b^{\ell+1}(Y) = (V_b^\ell(Y)) \setminus Y_{i(x)} = (V_b^\ell(Y)) \setminus \{x\} \subseteq (V_b^\ell(Y')) \setminus \{x'\} = (V_b^\ell(Y')) \setminus Y'_{i(x')} = V_b^{\ell+1}(Y'). \square$$

The claim implies that

$$\left(\max_{\Pi^{s_b^\ell}} V_b^\ell(Y')\right) \Gamma^{s_b^\ell} \left(\max_{\Pi^{s_b^\ell}} V_b^\ell(Y)\right) \quad (6)$$

for all ℓ ; this shows the result. \square

To see that Proposition 1 follows from Lemma B.1, we suppose that (1) holds for some $z, z' \in X$ and $Y \subseteq X$, and note that (1) also implies that

$$|Y \cup \{z\}| = |i(Y \cup \{z\})|. \quad (7)$$

Now, given (7), we know that if $z \notin C^b(Y \cup \{z\})$, then for each $s \in S_b$, the contract y assigned to s in the computation of $C^b(Y \cup \{z\})$ must be higher-priority than z under Π^s , that is, $y \Pi^s z$. But then, it follows from (7) and Lemma B.1 that each such s must be assigned a contract y' for which

$$y' \Gamma^s y \Pi^s z$$

in the computation of $C^b(Y \cup \{z, z'\})$. Thus, we must have $z \notin C^b(Y \cup \{z, z'\})$. Hence, we see that each C^b is substitutable on no-negotiation contract sets.

Proof of Proposition 2

See Appendix D.

Proof of Lemma 1

It is immediate that if $\varpi(\tilde{Y})$ is not individually rational, then \tilde{Y} is not individually rational for agents and slots. Thus, we need only consider the blocking conditions.

For $\tilde{Y} \subseteq \tilde{X}$, suppose that $Z \subseteq X$ is a set of contracts which blocks $\varpi(\tilde{Y})$. We fix some $b \in b(Z)$, and observe that there must be a contract $z \in Z_b \setminus \varpi(\tilde{Y})$ for which there is some step ℓ of the computation of $C^b(\varpi(\tilde{Y}) \cup Z)$ such that $D_b^\ell \setminus D_b^{\ell-1} = \{z\}$. (That is, there

must exist a contract $z \in Z_b \setminus \varpi(\tilde{Y})$ which is assigned to the highest-precedence slot, s_b^ℓ , among those slots which are assigned contracts in $Z_b \setminus \varpi(\tilde{Y})$ during the computation of $C^b(\varpi(\tilde{Y}) \cup Z)$.) We let $x \in \varpi(\tilde{Y})$ be the (possibly null) contract which is assigned to slot s_b^ℓ in the computation of $C^b(\varpi(\tilde{Y}))$.

It is clear that $z \Pi^{s_b^\ell} x$, by construction. Thus, we have $\langle z; s_b^\ell \rangle \tilde{\Pi}^{s_b^\ell} \langle x; s_b^\ell \rangle$. Meanwhile, we know that $z P^{i(z)}(\varpi(\tilde{Y}))_{i(z)}$ because Z blocks $\varpi(\tilde{Y})$. It follows that $\langle z; s_b^\ell \rangle$ is a slot-block for \tilde{Y} . Thus, if \tilde{Y} is not blocked at any slot, then $\varpi(\tilde{Y})$ is unblocked; the result follows.

Proof of Theorem 1

We prove the following result, which is slightly more general than Theorem 1.

Theorem B.1. *For any proposal order \sqsupset , the slot-outcome of the agent-optimal slot-stable mechanism in the agent-slot matching market corresponds (under projection ϖ) to the outcome of the cumulative offer process associated to proposal order \sqsupset .*

Proof. We suppress the dependence on \sqsupset , as doing so will not introduce confusion.

We begin with a simple lemma which shows that slots' assigned contracts improve (with respect to slot priorities) over the course of the cumulative offer process.

Lemma B.2. *Fix ℓ and ℓ' with $\ell < \ell'$, and let x^ℓ and $x^{\ell'}$ (with $b(x^\ell) = b = b(x^{\ell'})$) be the contracts assigned to $s \in S_b$ in the computations of $C^b(A_b^{\ell+1}) = D_b^\ell$ and $C^b(A_b^{\ell'+1}) = D_b^{\ell'}$, respectively. Then, $x^{\ell'} \Gamma^s x^\ell$.*

Proof. The result follows immediately from Lemma B.1, as $A_b^{\ell+1} \subseteq A_b^{\ell'+1} \subseteq X_b$ by construction. □

We denote the outcome of the cumulative offer process by Y , and let

$$\tilde{Y} \equiv \{\langle y; s \rangle : y \in Y \text{ and } s \text{ is assigned } z \text{ in the computation of } C^{b(z)}(Y)\}.$$

By construction, we have $\varpi(\tilde{Y}) = Y$.

Lemma B.3. *The slot-outcome \tilde{Y} is slot-stable.*

Proof. We suppose that $\langle z; s \rangle$ slot-blocks \tilde{Y} , so that

$$zP^{i(z)}(\varpi(\tilde{Y}))_{i(z)} = Y_{i(z)}, \quad (8)$$

$$z\Pi^s(\varpi(\tilde{Y}))_s = Y_s. \quad (9)$$

Now, by (8) and the fact that Y is the cumulative offer process outcome, we know that z must be proposed in some step $\underline{\ell}$ of the cumulative offer process. We let $\ell \geq \underline{\ell}$ be the first step of the cumulative offer process for which no slot $s' \in S_{b(z)}$ with $s' \triangleright^{b(z)} s$ is assigned z in the computation of $C^{b(z)}(A_{b(z)}^{\ell+1}) = D_{b(z)}^\ell$. (Such a step ℓ must exist since $z \notin Y$.) We let x^ℓ be the contract assigned to s in the computation of $C^{b(z)}(A_{b(z)}^{\ell+1})$. Since $x^\ell \neq z$, we know that $x^\ell \Pi^s z$. But then, we know by Lemma B.2 that for each $\ell' \geq \ell$, the contract $x^{\ell'}$ assigned to s in the computation of $C^{b(z)}(A_{b(z)}^{\ell'+1})$ has (weakly) higher Π^s -priority than x^ℓ , and hence has (strictly) higher Π^s -priority than z : $x^{\ell'} \Gamma^s x^\ell \Pi^s z$. In particular, then, we must have $Y_{i(z)} \Pi^s z$, contradicting (9). Thus, there cannot be a slot-contract $\langle z; s \rangle$ which slot-blocks \tilde{Y} . \square

Now, we let \tilde{Z} be the agent-optimal slot-stable slot-outcome. (Such an outcome exists by Theorem 3 of Hatfield and Milgrom (2005).)

Lemma B.4. *For each agent $i \in I$, $\tilde{Y}_i \tilde{R}^i \tilde{Z}_i$.*

Proof. It suffices to show that no contract $z \in \varpi(\tilde{Z})$ is ever rejected during the cumulative offer process. To see this, we suppose the contrary, and consider the first step ℓ at which some contract $z \in \varpi(\tilde{Z})$ is rejected. We let $s \in S_{b(z)}$ be the slot such that $\langle z; s \rangle \in \tilde{Z}$, and let $x \neq z$ be the contract assigned to s in the computation of $C^{b(z)}(A_{b(z)}^{\ell+1})$.

Now, as z is the first contract in $\varpi(\tilde{Z})$ to be rejected, we know that $xP^{i(x)}(\varpi(\tilde{Z}))_{i(x)}$. Moreover, as $z \notin C^{b(z)}(A_{b(z)}^{\ell+1})$ and x is assigned to s in the computation of $C^{b(z)}(A_{b(z)}^{\ell+1})$, we know that $x \Pi^s z$. But then, it follows that $\langle x; s \rangle$ slot-blocks \tilde{Z} , contradicting the fact that \tilde{Z} is slot-stable. \square

Now, by Lemma B.3, we know that \tilde{Y} is slot-stable; it then follows from Lemma B.4 that \tilde{Y} must be the agent-optimal slot-stable slot-outcome. The theorem then follows directly, as $\varpi(\tilde{Y}) = Y$. \square

Proof of Theorem 2

The result follows immediately from Lemma 1 and Theorem 1, as the agent-proposing deferred acceptance algorithm in the agent–slot matching market yields the agent-optimal slot-outcome (by Theorem 3 of Hatfield and Milgrom (2005)).

Proof of Theorem 3

We show Theorem 3 by way of the following more general result.

Theorem B.2. *For any stable outcome Y which is not equal to the outcome Z of the cumulative offer process, there is some agent $i \in I$ such that $Z_i P^i Y_i$.*

Proof. We suppose to the contrary that there is some stable outcome $Y \neq Z$ such that $Y_i R^i Z_i$ for all $i \in I$. We prove the theorem by considering an alternative proposal order \sqsupset' for the cumulative offer process, and showing that the outcome Z' of the cumulative offer process associated to \sqsupset' must be (weakly) preferred to Y by all agents.

We now describe the alternative proposal order:⁴¹ Let \sqsupset' be any order such that at each step of the cumulative offer process associated to \sqsupset' ,

1. all agents in $i(Y)$ who wish to propose contracts weakly preferred to those in Y have the opportunity to propose contracts before any agents in $I \setminus i(Y)$ do, and
2. all agents in $i(Y)$ who wish to propose contracts *not* weakly preferred to those in Y are not allowed to propose unless no agents in $I \setminus i(Y)$ wish to propose contracts.

Claim. *We have $Z'_i R^i Y_i$ for all $i \in I$.*

⁴¹For ease of exposition, we do not provide a fully formal specification of \sqsupset' ; nonetheless, our description directly yields an algorithm to compute such a specification.

Proof. We observe that if there is some step of the cumulative offer process associated to \sqsupset' at which all contracts in Y are held, then all future proposals will be rejected, and so the process produces the outcome $Z' = Y$.⁴² To see this, we observe that no agent in $i(Y)$ will propose while Y is held. Meanwhile, if at that time some agent $i \in I \setminus i(Y)$ proposes a contract x which is held following its proposal, then $\{x\}$ blocks Y , contradicting the fact that Y is stable.

Now, we note that by our choice of \sqsupset' , at any step in the cumulative offer process associated to \sqsupset' when there exists at least one agent $i \in i(Y)$ who wishes to propose a contract weakly preferred to Y_i , some such agent is given the opportunity to propose. It follows that, if there is no step of the cumulative offer process associated to \sqsupset' at which all contracts in Y are held, then the process must terminate *before* all contracts in Y have been proposed. In this case, the process outcome Z' must be weakly-preferred to Y by all agents; combining this fact with our previous observations shows the claim. \square

Now, by Theorem B.1, Z' must be the outcome of the cumulative offer process associated to the original proposal order—that is, $Z' = Z$. But, by the above claim, we then have

$$Z_i = Z'_i R^i Y_i R^i Z_i$$

for all $i \in I$, contradicting the assumption that $Z \neq Y$. \square

Theorem 2 shows that the outcome Z of the cumulative offer process is stable. But then, if there were an agent-optimal stable outcome $Y \neq Z$, we would have a stable outcome Y for which $Y_i R^i Z_i$ for all $i \in I$, contradicting Theorem B.2.

⁴²Formally, our argument shows that an analogous statement is true under any proposal order; we state the claim for proposal order \sqsupset' because that statement is important in the sequel.

Proof of Theorem 4

Part 1 is immediate from Theorem 2. Meanwhile, for Part 2, we suppose that there is some set of agents $I' \subset I$ and $\bar{P}^{I'} \neq P^{I'}$ such that

$$\varphi(\bar{P}^{I'}, P^{-I'}) P^i \varphi(P^{I'}, P^{-I'}) \quad (10)$$

for all $i \in I'$. Now, if \tilde{Z} is the outcome of the agent-optimal slot-stable mechanism run on preferences $(\tilde{\bar{P}}^{I'}, \tilde{P}^{-I'})$ and \tilde{Y} is the outcome of the agent-optimal slot-stable mechanism run on preferences $(\tilde{P}^{I'}, \tilde{P}^{-I'})$, then we have $\tilde{Z} \tilde{P}^i \tilde{Y}$ for all $i \in I'$, as we have

$$\varpi(\tilde{Z}) = \varphi(\bar{P}^{I'}, P^{-I'}) P^i \varphi(P^{I'}, P^{-I'}) = \varpi(\tilde{Y})$$

by (10) and Theorem B.1. But this implies that the agent-optimal slot-stable mechanism is not group strategy-proof (in the agent–slot market), contradicting Theorem 1 of Hatfield and Kojima (2009).⁴³

Proof of Theorem 5

To see this, we fix an agent i and let $\bar{\Pi}^B$ be an unambiguous improvement over Π^B for i . We let \sqsupset be the proposal order used in computing Φ , and let \sqsupset' be the alternative (uniquely defined) proposal order such that for all ℓ ,

$$\begin{aligned} j \sqsupset'_\ell k &\iff j \sqsupset_\ell k && \text{(for all } j, k \neq i) \\ j \sqsupset'_\ell i &&& \text{(for all } j \neq i), \end{aligned}$$

that is, the order obtained from \sqsupset by moving i to the bottom of each linear order \sqsupset_ℓ .

⁴³Theorem 1 of Hatfield and Kojima (2009) implies that the agent–slot matching mechanism which selects the (slot-)outcome of the agent-optimal slot-stable mechanism is group strategy-proof for agents. To see this, it suffices to note that both the substitutes condition and the law of aggregate demand hold automatically (for all preferences) in one-to-one matching markets such as the agent–slot matching market.

By Theorem B.1, the outcome $\Phi_{\bar{\Pi}}(P^I)$ is equal to that of the cumulative offer process associated to \sqsupset' under priorities $\bar{\Pi}$. Likewise, the outcome $\Phi_{\Pi}(P^I)$ is equal to that of the cumulative offer process associated to \sqsupset' under priorities Π . These observations essentially prove the result: In any cumulative offer process associated to \sqsupset' , agent i always proposes after all other agents' are unwilling to propose new contracts. Hence, under priority structure $\bar{\Pi}$, there is some contract x which i proposes in the last step before the cumulative offer process associated to \sqsupset' terminates. As $\bar{\Pi}$ is an unambiguous improvement over Π for i , we see that i proposes x in the cumulative offer process associated to \sqsupset' , under priorities Π .⁴⁴ It follows that

$$(\Phi_{\bar{\Pi}}(P^I))_i = xR^i(\Phi_{\Pi}(P^I))_i,$$

as desired.

Proof of Proposition 3

In any problem with agent types, $|X_i \cap X_b| = 1$ for all $i \in I$ and $b \in B$. It follows immediately that in such a problem,

$$|Y| = |i(Y)| \text{ for all } b \in B \text{ and } Y \subseteq X_b. \quad (11)$$

The substitutability of each branch choice function C^b in the presence of agent types then follows from Proposition 1, as substitutability on no-negotiation contract sets is equivalent to substitutability under condition (11). Additionally, we see that in the presence of agent types, each choice function C^b satisfies the law of aggregate demand, as condition (11) and Lemma B.1 together show that for all ℓ and $Y \subseteq Y'$, $\max_{\Pi^{s_b^\ell}} V_b^\ell(Y) = \emptyset_{s_b^\ell}$ whenever

⁴⁴As in the cumulative offer process associated to \sqsupset' , any contract x' with $i(x') = i$ and $x'P^i x$ is proposed before x is proposed. Moreover, by our choice of \sqsupset' , the process state at the time of the proposal of such a contract x' is exactly the same under priorities Π as it is under priorities $\bar{\Pi}$. Then, such x' must be rejected under priorities Π , as otherwise $\bar{\Pi}$ would not be an unambiguous improvement over Π for i .

$\max_{\Pi^{s_b^\ell}} V_b^\ell(Y') = \emptyset_{s_b^\ell}$; hence,

$$|C^b(Y')| = |C^b(Y')| \geq |C^b(Y)| = |C^b(Y)|.$$

Proof of Proposition 5

We let $b \in B$ be some branch for which $\triangleright^b \neq \blacktriangleright^b$, and let ℓ be the minimal value such that $s_b^\ell \in S_b^+$. As $\triangleright^b \neq \blacktriangleright^b$, there are $s \in S_b^0$ and $s' \in S_b^+$ such that $s' \triangleright^b s$. In particular, then, there must be some slot $s \in S_b^0$ for which $s_b^\ell \triangleright^b s$. We let ℓ' be such that $s_b^{\ell'}$ is the \triangleright^b -minimal such slot $s \in S_b^0$.

We label the cadets in I as i^1, i^2, \dots by the ranking π^* , so that

$$i^m \pi^* i^{m'} \iff m \leq m'.$$

We assume that

$$P^{i^m} = \begin{cases} (i^m, b, \mathbf{t}_0) \succ \emptyset_{i^m} & m < \ell \\ (i^m, b, \mathbf{t}_0) \succ (i^m, b, \mathbf{t}_+) \succ \emptyset_{i^m} & \ell \leq m \leq \ell' \\ (i^m, b, \mathbf{t}_0) \succ \emptyset_{i^m} & \ell' < m. \end{cases}$$

Claim. *The outcomes*

$$Y \equiv \{(i^m, b, \mathbf{t}) : m < \ell\} \cup \{(i^m, b, \mathbf{t}) : \ell \leq m < \ell'\} \cup \{(i^m, b, \mathbf{t}) : \ell' \leq m \leq |S_b|\},$$

$$Y' \equiv \{(i^m, b, \mathbf{t}) : m < \ell\} \cup \{(i^m, b, \mathbf{t}) : \ell < m \leq \ell'\} \cup \{(i^{\ell'}, b, \mathbf{t})\} \cup \{(i^m, b, \mathbf{t}) : \ell' < m \leq |S_b|\}$$

are both stable under the priorities Π and preferences P^I .

Proof. Clearly, both Y and Y' are individually rational. Thus, it suffices to show that each is unblocked.

Now under Y , all agents i^m for which $m < \ell$ or $m \geq \ell'$ hold their most-preferred contract. Thus, any set blocking Y must be a subset of $Z^* \equiv \{(i^m, b, \mathbf{t}_0) : \ell \leq m < \ell'\}$. However, the

contracts (i^m, b, t_+) are assigned to slots s_b^m in the computation of $C^b(Y)$ (for $m \leq |S_b|$), and $\tau^{s_b^m} = \tau^+$ for all $\ell \leq m < \ell'$. It follows that $Y = C^b(Y \cup Z)$ for any $Z \subseteq Z^*$; hence Y is unblocked, as desired. An analogous argument shows that Y' is unblocked. \square

The result follows directly from the claim, since $Y_{i\ell'} P^{i\ell'} Y'_{i\ell'}$ but $Y'_{i\ell} P^{i\ell} Y_{i\ell}$.

C Example Omitted from the Main Text

Example C.1. Let $X = \{i_b, i_{b'}, i'_b, i'_{b'}, i''_b, i''_{b'}, j_b, j_{b'}\}$, with $B = \{b, b'\}$, $I = \{i, i', i'', j\}$, $i(h_b) = h = i(h_{b'})$ for each $h \in I$, and $b(h_{b'}) = b''$ for each $h \in I$ and $b'' \in B$. We suppose that there are two types of agents— $\Gamma = \{i, j\}$ —and that $t(i) = t(i') = t(i'') = i$, while $t(j) = j$. We suppose that agents have preferences

$$P^i : i_b \succ i_{b'} \succ \emptyset_i$$

$$P^{i'} : i'_b \succ \emptyset_{i'}$$

$$P^{i''} : i''_{b'} \succ \emptyset_{i''}$$

$$P^j : j_b \succ j_{b'} \succ \emptyset_j.$$

We suppose further that b has two slots, $s_b^1 \triangleright^b s_b^2$, with priorities given by

$$\Pi^{s_b^1} : i_b \succ i'_b \succ i''_b \succ j_b \succ \emptyset_{s_b^1},$$

$$\Pi^{s_b^2} : i_b \succ j_b \succ i'_b \succ i''_b \succ \emptyset_{s_b^2},$$

and that b' has one slot, $s_{b'}^1$, with priority order identical to $\Pi^{s_b^2}$,

$$\Pi^{s_{b'}^1} : i_b \succ j_b \succ i'_b \succ i''_b \succ \emptyset_{s_{b'}^1}.$$

In this example, the cumulative offer process outcome is $\{i_b, j_b, i''_{b'}\}$. If the precedence of slots s_b^1 and s_b^2 were reversed, however, the cumulative offer process outcome would be

$\{i_b, i'_b, j_{b'}\}$; hence i'' is made worse off following a decrease in the precedence of a slot that favors agents of type $i = t(i'')$.

D Proof of Proposition 2

We first prove a lemma which shows that branch choice functions satisfy the *irrelevance of rejected contracts* property of Aygün and Sönmez (2012b,a).

Lemma D.1. *If $z' \notin C^b(Y \cup \{z, z'\})$, for some $z, z' \in X$ and $Y \subseteq X$, then $C^b(Y \cup \{z, z'\}) = C^b(Y \cup \{z\})$.*

Proof. We suppose that $z' \notin C^b(Y \cup \{z, z'\})$, and show the following claim.

Claim. *Let $V_b^\ell(Z)$ denote the set V_b^ℓ defined in step $\ell - 1$ of the computation of $C^b(Z)$. Then, $V_b^\ell(Y \cup \{z, z'\}) = V_b^\ell(Y \cup \{z\}) \cup \{z'\}$.*

Proof. We proceed by induction. We have $V_b^1(Y \cup \{z, z'\}) = V_b^1(Y \cup \{z\}) \cup \{z'\}$ a priori, so we assume that $V_b^{\ell'}(Y \cup \{z, z'\}) = V_b^{\ell'}(Y \cup \{z\}) \cup \{z'\}$ for all $\ell' < \ell + 1$ for some $\ell > 0$. We now show that this hypothesis implies that $V_b^{\ell+1}(Y \cup \{z, z'\}) = V_b^{\ell+1}(Y \cup \{z\}) \cup \{z'\}$: Let $x' \equiv \max_{\Pi^{s_b^\ell}}(V_b^\ell(Y \cup \{z, z'\}))$. As $z' \notin C^b(Y \cup \{z, z'\})$, we must have $x' \neq z'$; hence,

$$x' = \max_{\Pi^{s_b^\ell}}(V_b^\ell(Y \cup \{z, z'\}) \setminus \{z'\}) = \max_{\Pi^{s_b^\ell}}(V_b^\ell(Y \cup \{z\})) \quad (12)$$

by the inductive hypothesis. It then follows that

$$\begin{aligned} V_b^{\ell+1}(Y \cup \{z, z'\}) &= V_b^\ell(Y \cup \{z, z'\}) \setminus \{x'\} \\ &= (V_b^\ell(Y \cup \{z\}) \cup \{z'\}) \setminus \{x'\} \\ &= (V_b^\ell(Y \cup \{z\}) \setminus \{x'\}) \cup \{z'\} \\ &= V_b^{\ell+1}(Y \cup \{z\}) \cup \{z'\}, \end{aligned}$$

where the second equality follows from the inductive hypothesis, and the fourth equality follows from (12). This completes our induction. \square

The claim implies the desired result, as it shows that

$$\begin{aligned}
C^b(Y \cup \{z, z'\}) &= (Y \cup \{z, z'\}) \setminus V_b^{qb}(Y \cup \{z, z'\}) \\
&= (Y \cup \{z, z'\}) \setminus (V_b^{qb}(Y \cup \{z\}) \cup \{z'\}) \\
&= (Y \cup \{z\}) \setminus (V_b^{qb}(Y \cup \{z\})) = C^b(Y \cup \{z\}). \quad \square
\end{aligned}$$

Now, we suppose that $i(z), i(z') \notin i(Y)$, and that $z \notin C^b(Y \cup \{z\})$. Supposing that $z \in C^b(Y \cup \{z, z'\})$, we see by Lemma D.1 that $z' \in C^b(Y \cup \{z, z'\})$. This implies immediately that $i(z) \neq i(z')$, as no branch ever selects two contracts with the same agent.

Claim. *If z^ℓ and y^ℓ are the contracts assigned to s_b^ℓ in the computation of $C^b(Y \cup \{z, z'\})$ and $C^b(Y \cup \{z\})$, respectively, then $z^\ell \Gamma^{s_b^\ell} y^\ell$.*

Proof. Let $H_b^\ell(Z)$ denote the set H_b^ℓ defined in step ℓ of the computation of $C^b(Z)$. We proceed by double induction: Clearly, either $z^1 = \max_{\Pi^{s_b^1}}(Y \cup \{z, z'\}) = \max_{\Pi^{s_b^1}}(Y \cup \{z\}) = y^1$ or $z^1 = \max_{\Pi^{s_b^1}}(Y \cup \{z, z'\}) = z'$, so $z^1 \Gamma^{s_b^1} y^1$ and $i(H_b^1(Y \cup \{z, z'\})) \subseteq i(H_b^1(Y \cup \{z\}) \cup \{z'\})$. Thus, we suppose that $z^{\ell'} \Gamma^{s_b^{\ell'}} y^{\ell'}$ and $i(H_b^{\ell'}(Y \cup \{z, z'\})) \subseteq i(H_b^{\ell'}(Y \cup \{z\}) \cup \{z'\})$ for all $\ell' < \ell$.

We have

$$\begin{aligned}
z^\ell &= \max_{\Pi^{s_b^\ell}}(V_b^\ell(Y \cup \{z, z'\})) = \max_{\Pi^{s_b^\ell}}\left(\left((Y \cup \{z, z'\}) \setminus \left(Y_{i(H_b^{\ell-1}(Y \cup \{z, z'\}))}\right)\right)\right), \\
y^\ell &= \max_{\Pi^{s_b^\ell}}(V_b^\ell(Y \cup \{z\})) = \max_{\Pi^{s_b^\ell}}\left(\left((Y \cup \{z\}) \setminus \left(Y_{i(H_b^{\ell-1}(Y \cup \{z\}))}\right)\right)\right).
\end{aligned}$$

Since $i(z') \notin i(Y)$ and $i(z) \neq i(z')$, we have

$$\begin{aligned}
\left((Y \cup \{z\}) \setminus \left(Y_{i(H_b^{\ell-1}(Y \cup \{z\}))} \right) \right) &= \left((Y \cup \{z\}) \setminus \left(Y_{i(H_b^{\ell-1}(Y \cup \{z\}) \cup \{z'\})} \right) \right) \\
&\subseteq \left((Y \cup \{z\}) \setminus \left(Y_{i(H_b^{\ell-1}(Y \cup \{z, z'\}))} \right) \right) \\
&\subseteq \left((Y \cup \{z, z'\}) \setminus \left(Y_{i(H_b^{\ell-1}(Y \cup \{z, z'\}))} \right) \right), \tag{13}
\end{aligned}$$

where the first inclusion follows from the hypothesis that $i(H_b^{\ell-1}(Y \cup \{z, z'\})) \subseteq i(H_b^{\ell-1}(Y \cup \{z\}) \cup \{z'\})$.

The inclusion (13) implies that $z^\ell \Pi^{s_b} y^\ell$. This observation completes the first part of the induction. Moreover, it quickly yields the second part. To see this, we observe that if $i(z^\ell) \notin i(H_b^{\ell-1}(Y \cup \{z\}))$, then either $z^\ell = z'$ or $z^\ell = y^\ell$, since $z^\ell \Pi^{s_b} y^\ell$ and $i(z') \notin i(Y \cup \{z\})$. In either case, we have $i(H_b^\ell(Y \cup \{z, z'\})) \subseteq i(H_b^\ell(Y \cup \{z\}) \cup \{z'\})$. And finally, if $i(z^\ell) \in i(H_b^{\ell-1}(Y \cup \{z\}))$, then

$$\begin{aligned}
i(H_b^\ell(Y \cup \{z, z'\})) &= (i(H_b^{\ell-1}(Y \cup \{z, z'\})) \cup \{i(z^\ell)\}) \\
&\subseteq (i(H_b^{\ell-1}(Y \cup \{z\}) \cup \{z'\}) \cup \{i(z^\ell)\}) \\
&= (i(H_b^{\ell-1}(Y \cup \{z\})) \cup \{i(z')\} \cup \{i(z^\ell)\}) \\
&= (i(H_b^{\ell-1}(Y \cup \{z\})) \cup \{i(z')\}) \\
&\subseteq (i(H_b^\ell(Y \cup \{z\})) \cup \{i(z')\}) = (i(H_b^\ell(Y \cup \{z\}) \cup \{z'\})),
\end{aligned}$$

so the induction is complete.⁴⁵ □

Now, if $z \notin C^b(Y \cup \{z\})$, we know that for each $s \in S_b$, the contract y assigned to s in the computation of $C^b(Y \cup \{z\})$ must be higher-priority than z under Π^s , that is, $y \Pi^s z$. The preceding claim then shows that each such s must be assigned a contract y' in the computation of $C^b(Y \cup \{z, z'\})$ for which $y' \Gamma^s y \Pi^s z$. Thus, we must have $z \notin C^b(Y \cup \{z, z'\})$, contradicting our supposition to the contrary.

⁴⁵Here, the first inclusion follows from the inductive hypothesis.