Designing for Diversity in Matching*

Scott Duke Kominers
Society of Fellows
Harvard University

Tayfun Sönmez
Department of Economics
Boston College

July 25, 2013

Abstract

To encourage diversity, schools often "reserve" some slots for students of specific types. Students only care about their school assignments and contractual terms like tuition, and hence are indifferent among slots within a school. Because these indifferences can be resolved in multiple ways, they present an opportunity for novel market design.

We introduce a two-sided, many-to-one matching with contracts model in which agents with unit demand match to branches, which may have multiple slots available to accept contracts. Each slot has its own linear priority order over contracts; a branch chooses contracts by filling its slots sequentially. We demonstrate that in these *matching markets with slot-specific priorities*, branches' choice functions may not satisfy the substitutability conditions typically crucial for matching with contracts. Despite this complication, we are able to show that stable outcomes exist in this framework and can be found by a cumulative offer mechanism that is strategy-proof and respects unambiguous improvements in priority. Our results provide insight into the design of transparent affirmative action matching mechanisms. Additionally, they enable us to introduce a new market design application—allocation of airline seat upgrades—and show the value of a seemingly *ad hoc* administrative decision in the United States Military Academy's branch-of-choice program.

JEL classification: C78, D63, D78.

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

*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 Samson Alva, Susan Athey, Gadi Barlevy, Marco Bassetto, Gary Becker, Eric Budish, Darrell Duffie, Steven Durlauf, Drew Fudenberg, Pingyang Gao, Edward Glaeser, Jason Hartline, John William Hatfield, Sam Hwang, Matthew Jackson, Sonia Jaffe, Fuhito Kojima, Ellen Dickstein Kominers, Jonathan Levin, Robert Lucas, Paul Milgrom, Roger Myerson, Yusuke Narita, Derek Neal, Alexandru Nichifor, Muriel Niederle, Michael Ostrovsky, Mallesh Pai, Philip Reny, Assaf Romm, Alvin Roth, Hugo Sonnenschein, Peter Troyan, Harald Uhlig, Utku Ünver, Alexander Westkamp, E. Glen Weyl, members of the "Inequality: Measurement, Interpretation, and Policy" Working Group, and numerous seminar audiences. Kominers gratefully acknowledges the support of National Science Foundation grant CCF-1216095, an AMS–Simons travel grant, and the Human Capital and Economic Opportunity Working Group sponsored by the Institute for New Economic Thinking, as well as the hospitality of the Harvard Program for Evolutionary Dynamics.

1 Introduction

Mechanisms based on the *agent-proposing deferred acceptance algorithm* of Gale and Shapley (1962) have been adopted widely in the design of centralized school choice programs.[1] Deferred acceptance, first proposed for school choice by Abdulkadiroğlu and Sönmez (2003), is popular in practice because it is

1. *stable*, guaranteeing that no student ever envies a student with lower priority, and

2. dominant-strategy incentive compatible—*strategy-proof*—"leveling the playing field" by eliminating gains to strategic sophistication (Abdulkadiroğlu et al. (2006); Pathak and Sönmez (2008)).[2]

Many school districts (e.g., Chicago, Boston, and New York City) are concerned with issues of student diversity and have thus embedded affirmative action systems into their school choice programs. However, rendering deferred acceptance compatible with affirmative action requires modification of the algorithm.

In this paper, we observe that diversity, financial aid, or other concerns often cause agents' priorities to vary across a given institution's slots. 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 theory of *many-to-one matching with contracts* (Kelso and Crawford (1982); Hatfield and Milgrom (2005)).[3]

[1] These 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 (2013)), and assignment of K–12 students to public schools in Denver in 2012. Perhaps most significantly, a version of deferred acceptance has been recently been adopted by all (more than 150) local authorities in England (Pathak and Sönmez (2013)).

[2] Strategy-proofness is also useful because it enables the collection of true preference data for planning purposes.

[3] Like our model, the Westkamp (forthcoming) model of *matching with complex constraints* allows priorities to vary across slots (see also Braun et al. (2012)). While Westkamp (forthcoming) allows more general forms of interaction across slots than we allow in the present work, Westkamp (forthcoming) does not allow the variation in match contract terms essential for applications like airline upgrade allocation (novel to this work) and cadet–branch matching (introduced by Sönmez and Switzer (2013) and Sönmez (2013)).

To make this case, we first illustrate the impacts of modified deferred acceptance implementations used in the Chicago and Boston school choice programs. Then, we develop a model of matching with slot-specific priorities, which embeds classical priority matching settings (e.g., Balinski and Sönmez (1999); Abdulkadiroğlu and Sönmez (2003)), models of affirmative action (e.g., Kojima (2012); Hafalir et al. (2013)), and the cadet–branch matching framework (Sönmez and Switzer (2013); Sönmez (2013)), as well as a new market design problem we introduce: airline seat upgrade allocation.[4]

We advocate for a specific implementation of the *cumulative offer mechanism*, which generalizes agent-proposing deferred acceptance (Hatfield and Milgrom (2005); Hatfield and Kojima (2010)). Previous priority matching models have relied on the existence of agent-optimal stable outcomes to guarantee that this mechanism is strategy-proof. In markets 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 surprising 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.[5]

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 agent-optimal stable outcomes. The existence of agent-optimal stable outcomes in our general model may depend on several factors, including the number of different contractual arrangements agents and institutions may have, and the *precedence order* according to which institutions prioritize individual slots above others.[6]

[4]The priority structure in our framework also generalizes the priority matching analog of *leader-follower responsive preferences* introduced in the study of matching markets with couples (Klaus and Klijn (2005); Hatfield and Kojima (2010)).

[5]These conclusions allow us to re-derive several main results of the innovative Hafalir et al. (2013) approach to welfare-enhancing affirmative action.

[6]As we show in an application of our model (Proposition 6), 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. This is important because an agent-optimal stable mechanism removes all

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)).[7] 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 (2013); Sönmez (2013)). Nevertheless, the priority structure in our model gives rise to 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 Hatfield and Milgrom (2005) 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 outcome[8]; 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. These relationships are key to our main results.

Finally, we note that the generality of our framework enables novel market design applications. We present one such application as an example: the design of mechanisms for the allocation of airline seat upgrades. In this setting, customers have preferences over upgrade acquisision channels—elite status, cash, and reward points—and airline seating classes have slot-specific priorities. Because there are multiple mediums of exchange, the airlines' choice functions in general fail not only the substitutability condition but also the milder unilateral substitutability condition. While these failures of substitutability place airline seat upgrade allocation outside the reach of the prior literature, our results show that upgrade allocation can indeed be conducted through matching with contracts in a manner that is fair,

conflict of interest among agents in the context of stable assignment.

[7]Thus, 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)).

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

strategy-proof, and (unambiguous-)improvement-respecting.

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 introduce our airline seat upgrade allocation application, revisit the affirmative action applications, and briefly discuss applications to cadet–branch matching. Section 6 concludes. Most proofs are presented in the Appendix.

2 Motivating Applications

The mechanics of producing the assignment of students to school seats received little attention in the debates over *school choice* 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 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).[9]

Deferred acceptance 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 the proponents of school choice and the proponents of alternative, "neighborhood assignment" systems based

[9]This mechanism is often called the *student-optimal stable mechanism*.

on students' home addresses.

However, deferred acceptance 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 school choice 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 lead to subtly different priority matching mechanisms.

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

2.1 Affirmative Action at Chicago's Selective High Schools

Since 2009, Chicago Public Schools (CPS) has admitted students to its selective enrollment high schools using an affirmative action plan based on socio-economic status (SES) and a student allocation mechanism based on deferred acceptance.[10] Under the current plan for assigning students to schools, the SES of each student is determined based on home address; students are then divided into four roughly evenly sized tiers:[11] In 2009,

- 135,716 students living in 210 Tier 1 (lowest-SES) tracts had a median family income of $30,791,

[10]CPS initially adopted a version of the Boston mechanism, but immediately abandoned it in favor of deferred acceptance. Pathak and Sönmez (2013) have presented a detailed account of this midstream reform.

[11]SES 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.

- 136,073 students living in 203 Tier $\mathsf{2}$ tracts had a median family income of \$41,038,
- 136,378 students living in 226 Tier $\mathsf{3}$ tracts had a median family income of \$54,232, and
- 136,275 students living in 235 Tier $\mathsf{4}$ (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*.[12] 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 inaccessible to children from poorer neighborhoods. In order to implement this objective, CPS adopted the following priority structure at each of the nine selective enrollment high schools for the 2010–2011 school year:

- priority for 40% of the seats is determined by students' composite scores, while
- for each of the four SES tiers t, 15% of the seats are set aside for Tier t students, with composite score determining relative priority among those students.

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 breaks each selective enrollment high school into five hypothetical schools: The set S_b of seats at each school b is partitioned into subsets

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

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 t—they give Tier t students

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

priority over other students, and use composite scores to rank students within Tier $\mathtt{t}$. Each set of seats is viewed as a separate "school" within the CPS algorithm. Because seat priorities are uniform within each set $S_b^{\mathtt{e}}$, 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 $S_b^{\mathtt{e}}$ by $\langle i; s_b^{\mathtt{e}} \rangle$, then any Tier $\mathtt{t}$ student i is indifferent among

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

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

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

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

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

transformation used in the CPS implementation specifically aids the students whose scores are *low*, relative to the (predominantly high-SES) students allocated open seats.

To understand this effect, 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 students of 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 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.

The size of the score distribution blocks—and hence the size of the effect of switching to the counterfactual mechanism—is an empirical question. Using data from the 2010-2011 Chicago school choice admissions program, we now show that our intuition is accurate and

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.

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

Table 1: Comparison of individual student outcomes.

that the magnitude of this effect is substantial.[14]

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.

Of the 4270 seats allocated at nine selective high schools, 771 of them change hands between the two scenarios—what at first appears to be a minor decision in the implementation of the algorithm 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.

[14]Our data set includes the 2010-2011 quotas for the nine CPS elite public high schools, as well as students' composite scores and submitted rank-order lists.

	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
TOTAL	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.

- 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 and Figure 2 show 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.[15]

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 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.

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

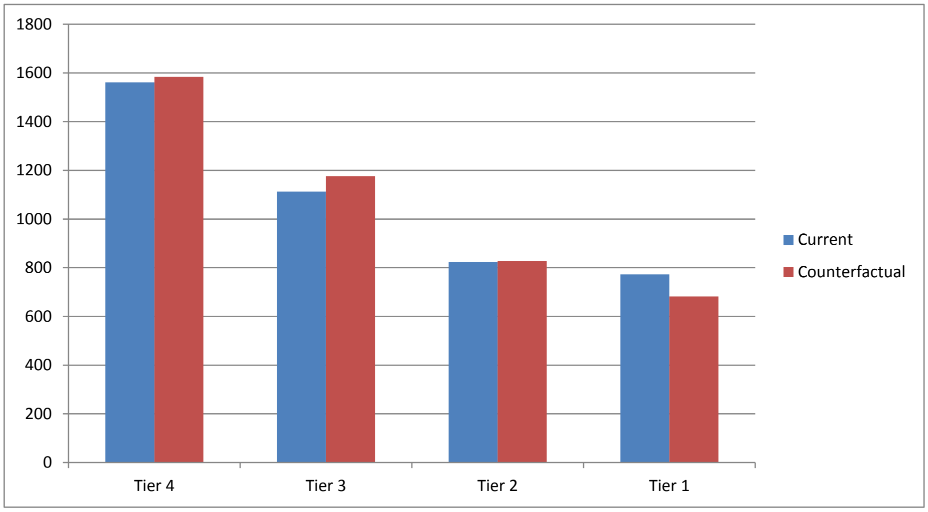

Figure 2: Effect of mechanism change on aggregate composition of selective high schools.

	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.

2.2 K–12 Admissions in Boston Public Schools

In the Boston school choice program, a student's priority 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 lottery number, which is used to break ties.

Half of the seats at each school (50%) give students with walk-zone priority higher claims, while the other half are "open," ranking students only on lottery number (and thus giving no special advantage to students with walk-zone priority). This choice of priority structure is not arbitrary—it represents a delicate balance between the interests of the school choice and neighborhood assignment advocates in Boston.

Thus, 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 (each having half of the true capacity)—one "walk-zone half" and one "open half." 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 walk-zone seats are ranked above open seats. Because of this implementation decision, the BPS school choice program systematically awards walk-zone seats to walk-zone students with lottery numbers high enough to acquire non-walk-zone seats, in effect having those students give up their high lottery draws. This effect vitiates Boston's walk-zone priority policy: As we show in a follow-up paper (Dur et al. (2013)), the outcome of Boston's implementation of the 50-50 seat split is nearly identical to the outcome under an alternative system in which there is no walk-zone priority at all (i.e. all seats are "open"). We show moreover that this is a direct consequence of Boston's choice of seat ordering—a reversal of the order, holding fixed the 50-50 split produces a significant increase in assignment of students to their walk-zone schools.

2.3 Uncertainty Regarding Student Performance

In practice, school choice programs' priority structures arise through a political bargaining process and may be adjusted to target specific outcomes. Indeed, CPS revised its priority structure slightly for the 2012–2013 school year.[16] (BPS, by contrast, has maintained its 50-50 split between walk-zone and open seats since it first implemented deferred acceptance, despite increased political emphasis on walk-zone assignment (Menino (2012)).)

Allowing for political bargaining, we might expect policymakers to adjust the number of reserved seats in reaction to the order in which seats are filled. Quota adjustment, however, in general cannot capture the full impact of the seat order on assignment outcomes. For instance, as the following stylized example illustrates, the choice of seat order affects the mechanism's behavior in the presence of uncertainty regarding the distribution of student test scores.[17]

Example 1. There are seven students, four of whom are of *majority* type, $\{\mathsf{M}_1, \mathsf{M}_2, \mathsf{M}_3, \mathsf{M}_4\}$, and three of whom are of *minority* type, $\{\mathsf{m}_1, \mathsf{m}_2, \mathsf{m}_3\}$. Each student i has a test score σ_i. The majority test score distribution is held fixed at $\sigma_{\mathsf{M}} = (\sigma_{\mathsf{M}_1}, \sigma_{\mathsf{M}_2}, \sigma_{\mathsf{M}_3}, \sigma_{\mathsf{M}_4}) \equiv (10, 9, 6, 5)$. With probability $\frac{1}{2}$, minorities score well, and their test score profile is $\sigma^{\mathrm{H}}_{\mathsf{m}} = (\sigma^{\mathrm{H}}_{\mathsf{m}_1}, \sigma^{\mathrm{H}}_{\mathsf{m}_2}, \sigma^{\mathrm{H}}_{\mathsf{m}_3}) \equiv (8, 7, 4)$, inducing rank order

$$\mathsf{M}_1 \succ \mathsf{M}_2 \succ \mathsf{m}_1 \succ \mathsf{m}_2 \succ \mathsf{M}_3 \succ \mathsf{M}_4 \succ \mathsf{m}_3.$$

With probability $\frac{1}{2}$, minorities score poorly, and have score profile $\sigma^{\mathrm{L}}_{\mathsf{m}} = (\sigma^{\mathrm{L}}_{\mathsf{m}_1}, \sigma^{\mathrm{L}}_{\mathsf{m}_2}, \sigma^{\mathrm{L}}_{\mathsf{m}_3}) \equiv (4, 3, 2)$, inducing rank-order

$$\mathsf{M}_1 \succ \mathsf{M}_2 \succ \mathsf{M}_3 \succ \mathsf{M}_4 \succ \mathsf{m}_1 \succ \mathsf{m}_2 \succ \mathsf{m}_3.$$

[16] In the new priority structure, 5% of the seats are reserved for hand-picking by principals, the fraction of open seats is reduced to 28.5%, and the fraction of reserved seats for each SES tier is increased to 16.625%.

[17] We thank Michael Ostrovsky for suggesting the structure of this example.

There is one school, with five seats available. Some of the seats are *open* seats (s^{o}), and simply prioritize students according to test scores. Other seats are *reserved* for minorities (s^{m}), ranking minority students above majority students and using test scores to break ties within groups.

First, we suppose that there are two reserved seats and (as in Boston) those seats are filled before the open seats. Under the induced priority structure—for either minority score profile realization—two minority students (m_1 and m_2) are assigned to the reserved seats, and majority students (M_1, M_2, and M_3) receive the open seats (see Figure 1).

$$\sigma = (\sigma_{\mathsf{M}}, \sigma_{\mathsf{m}}^{\mathrm{H}})$$

$$\begin{aligned}
s^{\mathsf{m}1} &: \boxed{\mathsf{m}_1} \succ \mathsf{m}_2 \succ \mathsf{m}_3 \succ \cdots \\
s^{\mathsf{m}2} &: \mathsf{m}_1 \succ \boxed{\mathsf{m}_2} \succ \mathsf{m}_3 \succ \cdots \\
s^{\mathsf{o}1} &: \boxed{\mathsf{M}_1} \succ \mathsf{M}_2 \succ \mathsf{m}_1 \succ \mathsf{m}_2 \succ \mathsf{M}_3 \succ \mathsf{M}_4 \succ \cdots \\
s^{\mathsf{o}2} &: \mathsf{M}_1 \succ \boxed{\mathsf{M}_2} \succ \mathsf{m}_1 \succ \mathsf{m}_2 \succ \mathsf{M}_3 \succ \mathsf{M}_4 \succ \cdots \\
s^{\mathsf{o}3} &: \mathsf{M}_1 \succ \mathsf{M}_2 \succ \mathsf{m}_1 \succ \mathsf{m}_2 \succ \boxed{\mathsf{M}_3} \succ \mathsf{M}_4 \succ \cdots
\end{aligned}$$

$$\sigma = (\sigma_{\mathsf{M}}, \sigma_{\mathsf{m}}^{\mathrm{L}})$$

$$\begin{aligned}
s^{\mathsf{m}1} &: \boxed{\mathsf{m}_1} \succ \mathsf{m}_2 \succ \mathsf{m}_3 \succ \cdots \\
s^{\mathsf{m}2} &: \mathsf{m}_1 \succ \boxed{\mathsf{m}_2} \succ \mathsf{m}_3 \succ \cdots \\
s^{\mathsf{o}1} &: \boxed{\mathsf{M}_1} \succ \mathsf{M}_2 \succ \mathsf{M}_3 \succ \mathsf{M}_4 \succ \mathsf{m}_1 \succ \mathsf{m}_2 \succ \cdots \\
s^{\mathsf{o}2} &: \mathsf{M}_1 \succ \boxed{\mathsf{M}_2} \succ \mathsf{M}_3 \succ \mathsf{M}_4 \succ \mathsf{m}_1 \succ \mathsf{m}_2 \succ \cdots \\
s^{\mathsf{o}3} &: \mathsf{M}_1 \succ \mathsf{M}_2 \succ \boxed{\mathsf{M}_3} \succ \mathsf{M}_4 \succ \mathsf{m}_1 \succ \mathsf{m}_2 \succ \cdots
\end{aligned}$$

Figure 1: Assignments under the two possible score realizations, in the case that the school fills two reserved seats before filling the open seats.

Now, we suppose instead that (as in Chicago) the reserved seats are filled after the open seats. If the school system seeks to assign the same expected number of seats to minorities, then it must convert one reserved seat into an open seat. Under the new priority structure, minorities receive two open seats and one reserved seat when $\sigma_{\mathsf{m}}^{\mathrm{H}}$ is realized, and receive only the reserved seat when $\sigma_{\mathsf{m}}^{\mathrm{L}}$ is realized (see Figure 2).

$$\sigma = (\sigma_{\mathsf{M}}, \sigma_{\mathsf{m}}^{\mathrm{H}})$$

$$\begin{aligned}
s^{\mathsf{o}1} &: \boxed{\mathsf{M}_1} \succ \mathsf{M}_2 \succ \mathsf{m}_1 \succ \mathsf{m}_2 \succ \mathsf{M}_3 \succ \mathsf{M}_4 \succ \cdots \\
s^{\mathsf{o}2} &: \mathsf{M}_1 \succ \boxed{\mathsf{M}_2} \succ \mathsf{m}_1 \succ \mathsf{m}_2 \succ \mathsf{M}_3 \succ \mathsf{M}_4 \succ \cdots \\
s^{\mathsf{o}3} &: \mathsf{M}_1 \succ \mathsf{M}_2 \succ \boxed{\mathsf{m}_1} \succ \mathsf{m}_2 \succ \mathsf{M}_3 \succ \mathsf{M}_4 \succ \cdots \\
s^{\mathsf{o}4} &: \mathsf{M}_1 \succ \mathsf{M}_2 \succ \mathsf{m}_1 \succ \boxed{\mathsf{m}_2} \succ \mathsf{M}_3 \succ \mathsf{M}_4 \succ \cdots \\
s^{\mathsf{m}1} &: \mathsf{m}_1 \succ \mathsf{m}_2 \succ \boxed{\mathsf{m}_3} \succ \mathsf{M}_1 \cdots
\end{aligned}$$

$$\sigma = (\sigma_{\mathsf{M}}, \sigma_{\mathsf{m}}^{\mathrm{L}})$$

$$\begin{aligned}
s^{\mathsf{o}1} &: \boxed{\mathsf{M}_1} \succ \mathsf{M}_2 \succ \mathsf{M}_3 \succ \mathsf{M}_4 \succ \mathsf{m}_1 \succ \mathsf{m}_2 \succ \cdots \\
s^{\mathsf{o}2} &: \mathsf{M}_1 \succ \boxed{\mathsf{M}_2} \succ \mathsf{M}_3 \succ \mathsf{M}_4 \succ \mathsf{m}_1 \succ \mathsf{m}_2 \succ \cdots \\
s^{\mathsf{o}3} &: \mathsf{M}_1 \succ \mathsf{M}_2 \succ \boxed{\mathsf{M}_3} \succ \mathsf{M}_4 \succ \mathsf{m}_1 \succ \mathsf{m}_2 \succ \cdots \\
s^{\mathsf{o}4} &: \mathsf{M}_1 \succ \mathsf{M}_2 \succ \mathsf{M}_3 \succ \boxed{\mathsf{M}_4} \succ \mathsf{m}_1 \succ \mathsf{m}_2 \succ \cdots \\
s^{\mathsf{m}1} &: \boxed{\mathsf{m}_1} \succ \mathsf{m}_2 \succ \mathsf{m}_3 \succ \mathsf{M}_1 \cdots
\end{aligned}$$

Figure 2: Assignments under the two possible score realizations, in the case that the school fills one reserved seat after filling the open seats.

We thus see that it is not possible to choose the number of reserved seats so as to always achieve the same outcomes using the reserves-at-bottom policy as under the reserves-at-top policy with two reserved seats. Even holding fixed the expected number of minorities admitted, the order in which slots are filled impacts the behavior of the mechanism significantly: in the reserves-at-top treatment, the variance of the number of minorities assigned seats is 0, while in the reserves-at-bottom treatment, it is

$$\left(\frac{1}{2}(3^2) + \frac{1}{2}(1^2)\right) - (2)^2 = 5 - 4 = 1.$$

The general model we provide in the remainder of this paper shows that *any* ordering of schools' seats gives rise to a priority matching mechanism that is stable, strategy-proof, and improvement-respecting. Thus, the order in which seats are filled—the *order of precedence*, in the sequel—can be adjusted flexibly according to policymakers' market design goals.

3 Matching with Slot-Specific Priorities

The applications described in Section 2 motivate a richer school choice model than those considered in the literature—one that accommodates slot-specific priorities at each school. School choice, however, is not the only application of matching theory that would benefit from such a generalization. Motivated by United States Army's recently introduced branch-for-service program, Sönmez and Switzer (2013) and Sönmez (2013) have introduced and analyzed the Army's *cadet-branch matching* problem, under which cadets can increase their priorities at the bottom 25% of slots of each Army branch, in exchange for "bidding" three additional years of service commitment.

Like in Chicago's and Boston's school choice programs, cadet–branch matching relies upon the possibility of priority variation across slots within branches. Unlike school choice, however, each cadet can match with branches under multiple contract *terms*. Thus, a fully general model must build on the richer matching with contracts framework (Kelso and Crawford (1982); Hatfield and Milgrom (2005)), in order to go beyond the basic structure of school choice and cover the Army's matching problem. In addition to extending the aforementioned market design applications, the generality of our model allows us to introduce a new market design application—airline seat upgrade allocation—that is beyond the scope of earlier models in the literature (see Section 5.1).

3.1 Agents, Branches, Contracts, and Slots

In a **matching problem with slot-specific priorities**, there is a set of **agents** I, a set of **branches** B, and a (finite) set of **contracts** X.[18] Each contract $x \in X$ is between an agent $\mathrm{i}(x) \in I$ and branch $\mathrm{b}(x) \in B$.[19] We extend the notations $\mathrm{i}(\cdot)$ and $\mathrm{b}(\cdot)$ to sets of contracts by setting $\mathrm{i}(Y) \equiv \cup_{y \in Y}\{\mathrm{i}(y)\}$ and $\mathrm{b}(Y) \equiv \cup_{y \in Y}\{\mathrm{b}(y)\}$. For $Y \subseteq X$, we denote $Y_i \equiv \{y \in Y : \mathrm{i}(y) = i\}$; analogously, we denote $Y_b \equiv \{y \in Y : \mathrm{b}(y) = b\}$.

Each agent $i \in I$ has a (linear) preference order P^i (with weak order R^i) over contracts in $X_i = \{x \in X : \mathrm{i}(x) = i\}$. 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.[20] We say that the contracts $x \in X_i$ for which $\emptyset_i P^i x$ are **unacceptable to** i.

Each branch $b \in B$ has a set S_b of slots; each slot can be assigned at most one contract in $X_b \equiv \{x \in X : \mathrm{b}(x) = b\}$. Slots $s \in S_b$ have (linear) priority orders Π^s (with weak orders Γ^s) over contracts in X_b. For convenience, we use the convention that $Y_s \equiv Y_b$ for $s \in S_b$. As with agents, we assume that each slot s ranks a "null contract" $\emptyset_s$ which represents remaining unassigned.[21] We set $S \equiv \cup_{b \in B} S_b$.

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

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

[19] A 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.

[20] We use the convention that $\emptyset_i P^i x$ if $x \in X \setminus X_i$.

[21] As with agents, we use the convention that $\emptyset_s \Pi^s x$ if $x \in X \setminus X_s$.

3.2 Choice and the Order of Precedence

For any agent $i \in I$ and $Y \subseteq X$, 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 $y \in Y_s$.[22]

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. We assume that for each $b \in B$, the slots in S_b are ordered according to a (linear) **order of precedence** $\rhd^b$. We denote $S_b \equiv \{s_b^1, \ldots, s_b^{q_b}\}$ with $q_b \equiv |S_b|$ and the understanding that $s_b^\ell \rhd^b s_b^{\ell+1}$ unless otherwise noted. The interpretation of $\rhd^b$ is that if $s_b \rhd^b s_b'$ then—whenever possible—branch b fills slot s_b before filling s_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 $Y \setminus Y_{\mathrm{i}(x^1)}$ of contracts in Y with agents other than $\mathrm{i}(x^1)$.
- This process continues in sequence, with each slot s_b^ℓ being assigned the contract x^ℓ which is $\Pi^{s_b^\ell}$-maximal among contracts in the set $Y \setminus Y_{\mathrm{i}(\{x^1,\ldots,x^{\ell-1}\})}$.

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

1. Let $H_b^0 \equiv \varnothing$, and let $V_b^1 \equiv Y$.

[22]Here, 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.

2. For each $\ell = 1, \ldots, q_b$:

    (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 $V_b^{\ell+1} = V_b^\ell \setminus Y_{\mathrm{i}(x^\ell)}$.

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 $C^b(Y)$.[23]

3.3 Stability

An **outcome** is a set of contracts $Y \subseteq X$. We restrict our attention throughout to ("feasible") outcomes that contain at most one contract for each agent, i.e. $Y \subseteq X$ for which $|Y \cap X_i| \leq 1$ for each $i \in I$.

We follow the Gale and Shapley (1962) tradition in focusing on match outcomes that 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 $i \in \mathrm{i}(Z)$.

[23] If no contract $x \in Y$ is assigned to slot $s_b^\ell \in S_b$ in the computation of $C^b(Y)$, then s_b^ℓ is assigned the null contract $\emptyset_{s_b^\ell}$.

3.4 Conditions on the Structure of Branch Choice

We now discuss the extent to which branch choice functions satisfy the conditions that have been key to previous analyses of matching with contracts models. For the most part, our observations are negative[24]; thus, they help contextualize our results and illustrate some of the technical difficulties that arise in our general framework.

3.4.1 Substitutability Conditions

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

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

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.[25]

Choice function substitutability is necessary (in the maximal domain sense) 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

[24] As we show, branch choice functions in general fail both the (Hatfield and Milgrom (2005)) substitutability and (Hatfield and Kojima (2010)) unilateral substitutability conditions, and need not satisfy the (Hatfield and Milgrom (2005)) law aggregate demand.

[25] The 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 (forthcoming)). 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 of Appendix D) stable outcomes, but also to guarantee that there is no conflict of interest among agents.[26]

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

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

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

Unilateral substitutability is a powerful condition; it has been applied in the study of cadet–branch matching mechanisms (Sönmez and Switzer (2013); Sönmez (2013)). 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 2. Let $X = \{i_1, i_2, j_2\}$, with $B = \{b\}$, $I = \{i, j\}$, $\mathrm{i}(i_1) = i = \mathrm{i}(i_2)$ and $\mathrm{i}(j_2) = j$.[27] If b has two slots, $s_b^1 \rhd^b s_b^2$, with priorities given by

$$\Pi^{s_b^1} : i_1 \succ \emptyset_{s_b^1},$$
$$\Pi^{s_b^2} : i_2 \succ j_2 \succ \emptyset_{s_b^2},$$

then C^b fails the unilateral substitutability condition: if we take $z = j_2$, $z' = i_1$, and $Y = \{i_2\}$, then $z = j_2 \notin C^b(\{i_2, j_2\}) = C^b(Y \cup \{z\})$, but $z = j_2 \in C^b(\{i_1, i_2, j_2\}) = C^b(Y \cup \{z, z'\})$, even though $\mathrm{i}(z) = \mathrm{i}(j_2) = j \notin \{i\} = \mathrm{i}(\{i_2\}) = \mathrm{i}(Y)$.

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

[26] As in the work of Hatfield and Milgrom (2005), an irrelevance of rejected contracts condition (which is naturally satisfied in our setting—see Footnote 25) is implicitly assumed throughout the work of Hatfield and Kojima (2010) (Aygün and Sönmez (2012)).

[27] Clearly (in order for $\mathrm{b}(\cdot)$ to be well-defined), we must have $\mathrm{b}(i_1) = \mathrm{b}(i_2) = \mathrm{b}(j_2) = b$, as $|B| = 1$.

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

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

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

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

This weak substitutability condition, first introduced by Hatfield and Kojima (2008), is in general necessary (in the maximal domain sense) for the guaranteed existence of stable outcomes (Hatfield and Kojima (2008), Proposition 1).

Proposition 1. *Every branch choice function C^b is weakly substitutable.*

The choice functions C^b also satisfy the slightly stronger bilateral substitutability condition introduced by Hatfield and Kojima (2010).

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

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

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

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 2, we present Proposition 2 only to illustrate the relationship between our work and the conditions introduced in the prior literature.[28]

[28]Combining 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 26) this logic implicitly

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).[29]

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

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

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

Example 3. Let $X = \{i_1, i_2, j_1\}$, with $B = \{b\}$, $I = \{i, j\}$, $\mathrm{i}(i_1) = i = \mathrm{i}(i_2)$ and $\mathrm{i}(j_1) = j$.[30] If b has two slots, $s_b^1 \rhd^b s_b^2$, with priorities given by

$$\Pi^{s_b^1} : i_1 \succ j_1 \succ \emptyset_{s_b^1},$$
$$\Pi^{s_b^2} : i_2 \succ \emptyset_{s_b^2},$$

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

$$|C^b(\{i_2, j_1\})| = |\{i_2, j_1\}| = 2 > 1 = |\{i_1\}| = |C^b(\{i_1, i_2, j_1\})|.$$

4 Main Theoretical Results

We now develop our general theoretical results: In Section 4.1, we associate our original market to a (one-to-one) matching market in which slots, rather than branches, compete for

requires an irrelevance of rejected contracts condition (Aygün and Sönmez (2012)). However, as Hatfield and Kojima (2010) pointed out, the bilateral substitutability 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.

[29] Alkan (2002) and Alkan and Gale (2003) introduced a related *cardinal monotonocity* condition.

[30] Clearly (in order for $\mathrm{b}(\cdot)$ to be well-defined), we must have $\mathrm{b}(i_1) = \mathrm{b}(i_2) = \mathrm{b}(j_1) = b$, as $|B| = 1$.

contracts. Next, in Section 4.2, we introduce the cumulative offer process and use properties of the agent–slot matching market to show that the cumulative offer process always identifies a stable outcome. We then show moreover, in Section 4.3, that the cumulative offer process selects the agent-optimal stable outcome if such an outcome exists and corresponds to a stable outcome in the agent–slot matching market. Finally, in Section 4.4, we show that the mechanism which selects the cumulative offer process outcome is stable, strategy-proof, and improvement-respecting.

4.1 Associated Agent–Slot Matching Market

To associate a (one-to-one) **agent–slot matching market** to our original market, we extend the contract set X to the set $\tilde{X}$ defined by

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

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

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

Meawhile, 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:

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

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

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

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 $\varpi : \tilde{X} \to X$ defined by

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

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

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

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 $\langle z; s\rangle = \tilde{C}^{\mathrm{i}(z)}(\tilde{Y} \cup \{\langle z; s\rangle\})$ and $\langle z; s\rangle = \tilde{C}^s(\tilde{Y} \cup \{\langle z; s\rangle\})$.

By construction, slot-stable outcomes project (under ϖ) to stable outcomes.

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 shows 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 $\ell = 1, 2, \ldots$:

Step 1. Some agent $i^1 \in I$ proposes his most-preferred contract, $x^1 \in X_{i^1}$. Branch $\mathrm{b}(x^1)$ *holds* x^1 if $x^1 \in C^{\mathrm{b}(x^1)}(\{x^1\})$, and *rejects* x^1 otherwise. Set $A^2_{\mathrm{b}(x^1)} = \{x^1\}$, and set $A^2_{b'} = \emptyset$ for each $b' \neq \mathrm{b}(x^1)$; 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 $\mathrm{b}(x^2)$ holds the contracts in $C^{\mathrm{b}(x^2)}(A^2_{\mathrm{b}(x^2)} \cup \{x^2\})$ and rejects all other contracts in $A^2_{\mathrm{b}(x^2)} \cup \{x^2\}$; branches $b' \neq \mathrm{b}(x^2)$ continue to hold all contracts they held at the end of Step 1. Set $A^3_{\mathrm{b}(x^2)} = A^2_{\mathrm{b}(x^2)} \cup \{x^2\}$, and set $A^3_{b'} = A^2_{b'}$ for each $b' \neq \mathrm{b}(x^2)$.

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

If at any time no agent is able to propose a new contract—that is, if all agents for whom no contracts are on hold have proposed all contracts they find acceptable—

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 when no agents wish to propose contracts.

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.[31] An analogous order-independence result is known for settings where priorities induce unilaterally substitutable branch choice functions (Hatfield and Kojima (2010)). However, as we illustrated in Example 2, slot-specific priorities may not induce unilaterally substitutable choice functions. Meanwhile, no order-independence result is known for the general class of bilaterally substitutable choice functions.[32]

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 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.[33] 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

[31] We make this statement concrete in Theorem B.1 of Appendix B.

[32] Although Hatfield and Kojima (2010) state their cumulative offer process algorithm without attention to the proposal order, they only prove independence of the proposal order in the case of unilaterally substitutable preferences.

[33] Here, 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 $C^b(A_b^{\ell+1})$.

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 Theorem 3 of Hatfield and Milgrom (2005). This implies that $\tilde{Y} = \tilde{Z}$, proving Theorem 1 as $\varpi(\tilde{Y}) = Y$.

Our agent–slot matching market construction is superficially similar to earlier work reducing many-to-one matching markets with responsive preferences to one-to-one matching markets.[34] This similarity is delusory, however. Our construction is used in showing existence and strategy-proofness, with most of the work done through the proof and application of Theorem 1, as just described. The goal is not, as in the work of Roth and Sotomayor (1989), to obtain an alternate characterization of the set of stable outcomes. Moreover, to the extent that we do provide a partial correspondence between stable sets through Lemma 1, that correspondence is far more delicate than the one obtained by Roth and Sotomayor (1989): it is not onto (see Example 4), and depends crucially on the explicit respect of slot-precedence in agents' extended preferences.

Theorem 1 implies that the cumulative offer process always terminates.[35] Moreover, it shows 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

[34] In the case of responsive preferences, Roth and Sotomayor (1989) showed that the stable outcomes in a many-to-one "college admissions" matching market (without contracts) correspond exactly to the set of stable outcomes in a one-to-one matching market obtained by treating the seats at each college as separate individuals (who share the college's responsive preference order).

[35] This fact can also be observed directly, as the set X is finite and the full set of contracts available, $\cup_{b \in B} A_b^{\ell}$, grows monotonically in ℓ.

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 $i \in I$. A stable outcome $Y \subseteq X$ which Pareto dominates all other stable outcomes is called an **agent-optimal stable outcome**.[36] For general slot-specific priorities, agent-optimal stable outcomes need not exist, as the following example shows.

Example 4. Let $X = \{i_0, i_1, j_0, j_1, k_0, k_1\}$, with $B = \{b\}$, $I = \{i, j, k\}$ and $\mathrm{i}(h_0) = h = \mathrm{i}(h_1)$ for each $h \in I$.[37] We suppose that $h_0 P^h h_1 P^h \emptyset_h$ for each $h \in I$, and that b has two slots, $s_b^1 \rhd^b s_b^2$, with slot priorities given by

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

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.

Even when agent-optimal stable outcomes *do* exist, the cumulative offer process may not select them. To see this, we consider a modification of Example 4.[38]

Example 5. Let $X = \{i_0, i_1, i_*, j_0, j_1, j_*, k_0, k_1\}$, with $B = \{b\}$, $I = \{i, j, k\}$ and $\mathrm{i}(h_0) = \mathrm{i}(h_1) = \mathrm{i}(h_*) = h$ for each $h \in \{i, j\}$ and $\mathrm{i}(k_0) = k = \mathrm{i}(k_1)$.[39] We suppose that $h_0 P^h h_* P^h h_1 P^h \emptyset_h$

[36] That 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 $Y' \neq Y$.

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

[38] We thank Fuhito Kojima for this example.

[39] Clearly (in order for $\mathrm{b}(\cdot)$ to be well-defined), we must have $\mathrm{b}(x) = b$ for each $x \in X$, as $|B| = 1$.

for each $h \in \{i, j\}$, that $k_0 P^k k_1 P^k \emptyset_k$, and that b has two slots, $s_b^1 \rhd^b s_b^2$, with slot priorities given by

$$\Pi^{s_b^1} : j_1 \succ k_1 \succ i_0 \succ j_0 \succ k_0 \succ \emptyset_{s_b^1},$$
$$\Pi^{s_b^2} : i_* \succ j_* \succ i_0 \succ i_1 \succ j_0 \succ j_1 \succ k_0 \succ k_1 \succ \emptyset_{s_b^2}.$$

In this setting, the outcome $Y \equiv \{j_1, i_0\}$ is the agent-optimal stable outcome. However, the outcome of the cumulative offer process is $Y'' \equiv \{j_1, i_*\}$, which is Pareto dominated by Y.

Although the cumulative offer process does not always select agent-optimal stable outcomes, in general, it does find agent-optimal stable outcomes when they correspond to slot-stable outcomes. This fact is a direct consequence of Theorem 2.

Theorem 3. *If an agent-optimal stable outcome exists and is the projection (under ϖ) of a slot-stable outcome, then it is the outcome of the cumulative offer process.*

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

4.4.1 Stability and Strategy-Proofness

A mechanism φ is **stable** if it always selects an outcome stable with respect to slot priorities and input preferences. This condition dates back to Gale and Shapley (1962) and has been the backbone of the matching market design literature. It captures the natural idea that a mechanism produces an outcome consistent with the policy objectives reflected in the priority structure.

We say that a mechanism φ is **strategy-proof** if truthful preference revelation is a dominant strategy for agents $i \in I$, i.e. there is no agent $i \in I$, preference profile $P^I \in \prod_{j\in I} \mathcal{P}^j$, and $\bar{P}^i \neq P^i$ such that $\varphi\big(\bar{P}^i, P^{-i}\big) P^i \varphi\big(P^i, P^{-i}\big)$. Similarly, we say that a mechanism φ is **group strategy-proof** if there is no set of agents $I' \subset I$, preference profile $P^I \in \prod_{j\in I} \mathcal{P}^j$, and $\bar{P}^{I'} \neq P^{I'}$ such that $\varphi\big(\bar{P}^{I'}, P^{-I'}\big) P^i \varphi\big(P^{I'}, P^{-I'}\big)$ for all $i \in I'$. Strategy-proofness conditions have been central to the recent revolution in school choice market design because they eliminate benefits of strategic sophistication and costly strategic behavior, and enable the collection of true preference data (Abdulkadiroğlu et al. (2006); Pathak and Sönmez (2008)).

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 $\varphi\big(\bar{P}^{I'}, P^{-I'}\big) P^i \varphi\big(P^{I'}, P^{-I'}\big)$ for all $i \in I'$ would give rise to a profitable manipulation $(\tilde{\bar{P}}^{I'} \neq \tilde{P}^{I'})$ 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 Respect for 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 $y \in (X_{I\setminus\{i\}} \cup \{\emptyset_s\})$, 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. We say that φ **respects unambiguous improvements** if it respects unambiguous improvements for each agent $i \in I$.

While present in the matching literature since the work of Balinski and Sönmez (1999), respect for unambiguous improvements has not been central to previous debates on real-world market design.[40] Nevertheless, respect for improvements is essential in settings like the airline seat upgrade application we introduce in Section 5.1—there, it implies that customers never want to decrease their frequent-flyer status levels.[41] Similarly, respect for unambiguous improvements is important in cadet–branch matching (see Section 5.3), where cadets can influence their priority rankings directly—and may (in the absence of respect for

[40]Respect for improvements is, however, of importance in the growing normative literature on school choice design. For example, Hatfield et al. (2012) have used this condition in analyzing how school choice mechanism selection can impact schools' incentives for self-improvement.

[41]Formally, for this claim we need the extremely natural assumption that an increase in the status of customer i, holding other customers' status levels fixed, results in an unambiguous improvement in i's priority.

improvements) take perverse steps to lower their priorities.[42]

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 respect of unambiguous improvements. Theorem 5 would then follow from Theorem 1.[43]

[42]As Sönmez (2013) has illustrated, the current ROTC cadet–branch matching mechanism rewards cadets who can lower their priorities to just below the 50-th percentile mark. Evidence from Service Academy Forums (2012) suggests that cadets have figured this out, and may be adjusting their training and academic performance accordingly:

> "20% in the complete OML [order of merit list] might actually be 28% in the 'Active Duty' OML, so make sure you make this mental conversion to the complete OML during your first three years. Or, just really screw up everything except for GPA, and get yourself into the 55% (from the top = 45%) where you get your choice of Branch... just kidding. But in all seriousness, why create a system of merit evaluation that takes a top 40% OML cadet and rewards him/her for purposely sabotaging things to go DOWN in the OML to below the 50% AD OML line[...]?"

[43]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

In this section, we present applications of our theoretical results to the design of airline seat upgrade mechanisms, affirmative action programs, and cadet–branch matching mechanisms.

5.1 Airline Seat Upgrades

As the demand for airline seat upgrades has increased, airlines have begun providing multiple channels for obtaining upgrades. The most common channels for upgrades are:

1. automatic upgrades through *elite status*,
2. upgrades purchased with *cash* payments, and
3. upgrades purchased with *reward points* (i.e. *miles*), possibly together with some cash payments.

Typically, not all available upgrades in a given seat class can be obtained through elite status or miles, as airline companies, under pressure to increase profits, often place a quota restriction on the number of "rewards-based" upgrades. Similarly, because of profit pressures, airline companies might prioritize upgrading through cash payment.

Currently there is no unified "market clearing" system that allows customers to pursue upgrades via multiple upgrade channels. Instead, customers often need to pick one of the three channels—status, cash, or miles—when they decide to seek an upgrade.

To make this point clear, we consider the at-the-gate upgrade process immediately before a flight. A customer who is interested in an upgrade typically approaches the flight desk and inquires whether an upgrade is available. If an upgrade is available, those interested pursue upgrades through their preferred channels. Airline company personnel then determine the assignment of upgrades, considering factors including the composition of the set of customers interested in upgrades, those customers' frequent-flyer statuses, and the company's imposed quota/reserve restrictions.

Now consider an elite status customer i whose first choice is a free automatic upgrade, but who is willing to buy a cash upgrade as his second choice if he cannot receive a free upgrade. Under the current (standard) procedure it is not possible for customer i to express his preferences. Indicating willingness to buy the cash upgrade is essentially equivalent to giving up the possibility of a free elite status upgrade. Likewise a customer whose first choice is a miles upgrade and whose second choice is a cash upgrade is forced to pick only one of these options.

The airline seat upgrade allocation problem can be modeled as an application of our model in which the slot-specific priority structure gives the airlines the ability to implement a wide variety of allocation policies. The cumulative offer mechanism in this context offers airlines a convenient market clearing mechanism that allows customers to express their full preferences over different seating classes and types of upgrade. In this application branches correspond to different *seating classes* (e.g., business class and first class) while a contract of a customer specifies a *seating class and an upgrade channel* (e.g., a business class seat obtained through a payment of miles).

This is a novel application of two-sided matching with contracts not covered by the earlier literature. To see this, we give a simple example in which the airline's choice function fails not only the substitutability condition but also the milder unilateral substitutability condition—the weakest condition in the literature that has enabled applications of matching with contracts thus far.[44]

Example 6. There is a unique branch b, the *business class*, which has two upgrade slots available. One of the slots, slot s_b^1, accepts only cash upgrades, while the second slot, s_b^2 accepts both miles and cash but prioritizes cash. For each slot, ties between customers are broken with respect to frequent flyer status, and between the two slots the airline first tries to fill the less-permissive slot s_b^1 and then tries to fill slot s_b^2; in our terminology slot s_b^1 has higher precedence than slot s_b^2.[45]

[44]Note that this example is closely analogous to Example 2.

[45]This way, if there is one cash-customer and one miles-customer, the cash-customer will not block the

Consider two customers, i and j, with i having higher frequent-flyer status than j, and the following three contracts: $\{i_\$, i_m, j_m\}$, where $i_\$$ represents the cash upgrade contract for customer i and i_m and j_m respectively represent the mile upgrade contracts for customers i and j. Given the airline upgrade policy as described, we have $s_b^1 \rhd^b s_b^2$, and

$$\Pi^{s_b^1} : i_\$ \succ \emptyset_{s_b^1},$$
$$\Pi^{s_b^2} : i_\$ \succ i_m \succ j_m \succ \emptyset_{s_b^2}.$$

Observe that the resulting business class choice function C^b fails the unilateral substitutability condition (introduced on page 20): For $z = j_m$, $z' = i_\$$, and $Y = \{i_m\}$, we have

$$z = j_m \notin \{i_m\} = C^b(\{i_m, j_m\}) = C^b(Y \cup \{z\}),$$

while we have

$$z = j_m \in \{i_\$, j_m\} = C^b(\{i_m, j_m, i_\$\}) = C^b(Y \cup \{z, z'\})$$

even though $\mathrm{i}(z) = \mathrm{i}(j_m) = j \notin \{i\} = \mathrm{i}(i_m) = \mathrm{i}(Y)$.

Our discussion of airline seat allocation through matching with slot-specific priorities also illustrates how our framework could be used to implement the *seat upgrade auctions* several major airlines have begun implementing in the last year (McCartney (2013)). In this context, customers' preferences could express preferences over potential "bids" within each upgrade channel (e.g., a bid of \$200 might be preferred to a bid of 15000 points, which in turn is preferred to a bid of \$300); the cumulative offer mechanism then corresponds to an ascending auction for upgrades. Once again, in this application the airline choice functions may fail the unilateral substitutability condition.

More generally, while we have presented the specific application of our framework to questions of airline seat upgrade allocation, the same approach can be used in any market

airline from assigning both upgrade seats.

where there are multiple mediums of exchange and slot-specific-priorities.[46]

5.2 Affirmative Action Mechanisms

We say that a matching problem has **agent types** if the contract set X is a subset of $I \times B \times \mathsf{T}$ for some **type set** T, and for each $i \in I$, $X_i = \{i\} \times B \times \{\mathsf{t}\}$ for some $\mathsf{t} \in \mathsf{T}$, so that each i is associated to exactly one **type** t.[47] For such a problem, we identify agents with their types, writing $\mathrm{t}(i)$ for the unique type $\mathsf{t} \in \mathsf{T}$ such that $X_i = \{i\} \times B \times \{\mathsf{t}\}$. For consistency with the prior literature on school choice, we abuse notation slighly by writing i to denote, for each branch $b \in B$, the unique contract $(i, b, \mathrm{t}(i)) \in (X_i \cap X_b)$.

Imposing agent type structure on our general model simplifies the behavior of branches' choice functions.

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 weak substitutability—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.

In the presence of agent types, the structure of the set of stable outcomes is also simplified: stability and slot-stability exactly correspond.

[46] One possible structure for such a market is discussed by Peranson (2005) in the context of medical matching: A hospital may have multiple positions to fill, and prefer a balance between clinical-specialty doctors and research-focused doctors (see Example 2 of Peranson (2005)). In this setting, a doctor capable of both clinical and research work can sign either a *clinical* contract or a *research* contract; the hospital's preferences can be expressed using slot-specific structure in which some slots prioritize clinical contracts over research and others prioritize research contracts over clinical. Hospital preferences in this application are analogous to those of airlines in the problem of upgrade allocation ("clinical" and "research" are assignment channels akin to "cash" and "miles"), and depending on the structure of hospital preferences, this application may also fail the unilateral substitutability condition. That said, because allocation of doctors to hospitals is a two-sided matching problem (rather than a pure question of allocation under priorities), this application entails strategic concerns not analyzed in our framework.

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

Proposition 4. *In a matching problem with slot-specific priorities and agent types, every stable outcome is the projection of a slot-stable outcome.*

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 Proposition 4 and 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 $\mathtt{t}$ highly can improve type-$\mathtt{t}$ agents' cumulative offer outcomes. Unfortunately, while we believe this comparative static should generally hold in reasonably-sized marketplaces, it may fail in small markets.[48]

5.2.1 "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. (2013) 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. (2013) 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. (2013) embeds into the framework of matching with slot-specific priorities and agent types as follows: The agents $i \in I$ are *students* and the branches $b \in B$ are

[48] Example C.1 of Appendix C illustrates this fact.

schools. Each student $i \in I$ as a strict linear preference order P^i over schools, and is of either minority (m) or majority (M) type (i.e. $\mathrm{t}(i) \in \{\mathsf{m}, \mathsf{M}\} = \mathsf{T}$). Each school $b \in B$ has a strict linear "tiebreaker" 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 $r_b^{\mathsf{m}} \leq q_b$ such that b prefers any minority applicant to any majority applicant if the number of minority students admitted is below r_b^{m} (Hafalir et al. (2013)). This policy can be implemented by choosing slot-specific priorities $\bar{\Pi}$ such that

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

Under **affirmative action with majority quotas**, meanwhile, each school $b \in B$ has an associated majority quota $q_b^{\mathsf{M}} \leq q_b$ 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 $\ell < q_b - q_b^{\mathsf{M}}$,
   (a) $i \Pi^{s_b^\ell} i' \Pi^{s_b^\ell} \emptyset_{s_b^\ell} \iff \mathrm{t}(i) = \mathrm{t}(i') = \mathsf{m}$ and $i \pi^b i'$, and
   (b) $\emptyset_{s_b^\ell} \Pi^{s_b^\ell} i \iff \mathrm{t}(i) = \mathsf{M}$;
2. for all $\ell \geq q_b - q_b^{\mathsf{M}}$, $i \Pi^{s_b^\ell} i' \Pi^{s_b^\ell} \emptyset_{s_b^\ell} \iff i \pi^b i'$.

With these observations, we may derive two of the main results of Hafalir et al. (2013) as consequences of our general results for slot-specific priority structures.

Proposition 5 (Hafalir et al. (2013))**.** *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 $r_b^{\mathsf{m}} = q_b - q_b^{\mathsf{M}}$ 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 tiebreaker order profile 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

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

for each $i \in I$. Now, if we have $Y = \Phi_{\Pi}(P^I) = \Phi_{\bar{\Pi}}(P^I) = Z$, 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. □

5.2.2 Socioeconomic Affirmative Action in Chicago School Choice

We now demonstrate that our framework embeds the real-world structure of the Chicago selective high school 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: $\mathsf{T} = \{4, 3, 2, 1\}$.

The top 40% of the q_b slots at shool $b \in B$ 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

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

Figure 3: 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 ($\{\mathsf{4}, \mathsf{3}, \mathsf{2}, \mathsf{1}\}$).

$\mathrm{t}(i) = \mathsf{4}$; the next 15% are reserved for students $i \in I$ of type $\mathrm{t}(i) = \mathsf{3}$; and so forth. These priority structures are illustrated in Figure 3.[49]

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

$$S_b = S_b^{\mathsf{o}} \cup S_b^{\mathsf{4}} \cup S_b^{\mathsf{3}} \cup S_b^{\mathsf{2}} \cup S_b^{\mathsf{1}},$$

with S_b^{o} consisting of $\frac{40}{100} q_b$ slots, and each set S_b^{t} consisting of $\frac{15}{100} q_b$ slots.[50] The priorities of slots $s \in S_b^{\mathsf{o}}$ are such that

$$i \Pi^s i' \Pi^s \emptyset_s \qquad \Longleftrightarrow \qquad i \pi^* i'.$$

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

$$i \Pi^s i' \Pi^s \emptyset_s \qquad \Longleftrightarrow \qquad i \pi^* i'$$

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

$$s^{\mathsf{o}} \rhd^b s^{\mathsf{4}} \rhd^b s^{\mathsf{3}} \rhd^b s^{\mathsf{2}} \rhd^b s^{\mathsf{1}} \tag{4}$$

for all $s^{\mathsf{o}} \in S_b^{\mathsf{o}}$, $s^{\mathsf{4}} \in S_b^{\mathsf{4}}$, $s^{\mathsf{3}} \in S_b^{\mathsf{3}}$, $s^{\mathsf{2}} \in S_b^{\mathsf{2}}$, and $s^{\mathsf{1}} \in S_b^{\mathsf{1}}$.[51]

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

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

[51] As all slots $s^{\mathsf{e}} \in S_b^{\mathsf{e}}$ have identical priorities, (4) suffices to specify the precedence order up to equivalence.

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

Figure 4: 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.

Tables 1–3 show the effect of switching from the current precedence orders $\rhd^b$ to the alternate orders, illustrated in Figure 4, in which open slots are filled after reserve slots. Formally, these counterfactual precedence orders $\blacktriangleright^b$ are such that

$$
s^{\mathsf{4}} \blacktriangleright^b s^{\mathsf{3}} \blacktriangleright^b s^{\mathsf{2}} \blacktriangleright^b s^{\mathsf{1}} \blacktriangleright^b s^{\circ}
$$

for all $s^{\circ} \in S^{\circ}_b$, $s^{\mathsf{4}} \in S^{\mathsf{4}}_b$, $s^{\mathsf{3}} \in S^{\mathsf{3}}_b$, $s^{\mathsf{2}} \in S^{\mathsf{2}}_b$, and $s^{\mathsf{1}} \in S^{\mathsf{1}}_b$.

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 open slots are filled last. 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.[52] This approach spreads the access to open slots evenly throughout the priority structure.

Simulation results presented in Table 4 show that student outcomes under the intermediate priorities are almost identical to those arising when all reserved slots are filled before the open slots (the counterfactual discussed in Section 2.1). While this might at first seem sur-

[52] When 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.

	Counterfactual Mechanism (Open Slots Last)				Effect of Switching (to Intermediate Mechanism)			
	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
TOTAL	1584	1176	828	682	−8	5	3	0

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

prising, it 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 slots, then the highest-scoring low-SES students receive the first reserved slots. Then, once again, the top of the truncated scoring distribution consists of high-SES students; these students take the next open slots. The highest-scoring low-SES students who remain unassigned then receive reserved slots, 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 slots are filled before the open slots.

5.3 Cadet–Branch Matching

The **cadet–branch matching problem** studied by Sönmez and Switzer (2013) and Sönmez (2013) is a slot-specific priority matching problem with contract set $X = I \times B \times \{\mathsf{t}_0, \mathsf{t}_+\}$. 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 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 $\mathsf{t}, \mathsf{t}' \in \{\mathsf{t}_0, \mathsf{t}_+\}$,

$$(i, b, \mathsf{t})\Pi^s(i', b, \mathsf{t}') \quad \Longleftrightarrow \quad i\pi^* i'.$$

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

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

for all slots $s \in S_b^0$. Slots $s \in S_b^+$ are **branch-of-choice ("bidding") slots**, which give priority to t_+ contracts: for any $b \in B$, $s \in S_b^+$, and $i \neq i' \in I$,

$$\begin{gathered}(i, b, \mathsf{t}_+)\Pi^s(i', b, \mathsf{t}_0), \quad \text{and} \\ (i, b, \mathsf{t})\Pi^s(i', b, \mathsf{t}) \quad \Longleftrightarrow \quad i\pi^* i'\end{gathered}$$

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

In the settings of Sönmez and Switzer (2013) and Sönmez (2013), 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

$$s \in S_b^0 \text{ and } s' \in S_b^+ \quad \Longrightarrow \quad s \blacktriangleright^b s'. \tag{5}$$

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 (2013) 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 (2013) to propose the use of a cadet-optimal stable mechanism for cadet–branch matching.[53]

Our next result shows that the structure found by Sönmez and Switzer (2013) 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 6. *For any cadet–branch matching problem precedence order* $\rhd \neq \blacktriangleright$ *for which there exists* $b \in B$, $s \in S_b^0$, *and* $s' \in S_b^+$ *such that* $s' \rhd^b s$, *there exists a profile of cadet preferences under which no outcome stable with respect to the branch choice functions* C^b $(b \in B)$ *induced by the slot priorities* Π^s $(s \in S)$ *and precedence order* $\rhd$ *is cadet-optimal.*

6 Conclusion

In this paper, we have studied *slot-precedence*, a feature of priority structure that is present in many real-world applications of matching theory but has not been treated formally in the prior literature. As we have shown, the choice of precedence order has important distributional implications, but does not affect agents' strategic incentives. Thus, attention to the precedence order provides market designers and policymakers an additional degree of freedom in the design of priority matching mechanisms. This additional flexibility is particularly useful for diversity-motivated designs, like affirmative action systems.

Our work also has theoretical implications: We have shown that the existence of agent-optimal stable outcomes is not necessary for strategy-proof stable matching, and have reinforced and expanded the matching with contracts framework. Additionally, our general model allows us to address novel market design applications like the allocation of airline seat upgrades, clarifies the relationship between existing models of affirmative action (Kojima (2012); Hafalir et al. (2013)), and illustrates special structure present in the Army's specific choice of priority structure for cadet–branch matching (Sönmez and Switzer (2013); Sönmez

[53] In his discussion of market design for the ROTC cadet–branch match, Sönmez (2013) extended these results to the case in which more than two distinct contract terms are available.

(2013)).

Our model is not the most comprehensive priority matching framework possible, and some of our substantive results may extend to more general settings. Nevertheless, our slot-specific priorities framework naturally embeds nearly all of the priority structures currently in application.[54] Our approach allows us to conduct comparative static exercises showing how slot precedence affects outcomes in real-world school choice programs. Precedence orders also induce attractive theoretical structure, which allows us to link our model to the simpler problem of one-to-one agent–slot matching.

References

Abdulkadiroğlu, Atila and Tayfun Sönmez, "School choice: A mechanism design approach," *American Economic Review*, 2003, *93*, 729–747.

_ , 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. Massachusetts Institute of Technology Working Paper.

_ , _ , 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, Hiroyuki, "On a characterization of stable matchings," *Economics Letters*, 2000, *68*, 43–49.

[54]The priority structure of which we are aware that is not covered by our model is that used in the German university admission system (Westkamp (forthcoming); Braun et al. (2012)).

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. Boston College Working Paper.

_ **and** _ , "Matching with contracts: Comment," forthcoming. American Economic Review.

Balinski, Michel and Tayfun Sönmez, "A tale of two mechanisms: Student placement," *Journal of Economic Theory*, 1999, *84*, 73–94.

Braun, Sebastian, Nadja Dwenger, Dorothea Kübler, and Alexander Westkamp, "Implementing quotas in university admissions: An experimental analysis," 2012. Max Planck Institute for Tax Law and Public Finance Working Paper.

Dur, Umut, Scott Duke Kominers, Parag A. Pathak, and Tayfun Sönmez, "The demise of walk zones in Boston: Priorities vs. precedence in school choice," 2013. Boston College Working Paper.

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, Tamás, "A fixed point approach to stable matchings and some applications," *Mathematics of Operations Research*, 2003, *28*, 103–126.

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*, 2013, *8*, 325–363.

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. Stanford Graduate School of Business Working Paper.

Kelso, Alexander S. and Vincent P. Crawford, "Job matching, coalition formation, and gross substitutes," *Econometrica*, 1982, *50*, 1483–1504.

Klaus, Bettina and Flip Klijn, "Stable matchings and preferences of couples," *Journal of Economic Theory*, 2005, *121*, 75–106.

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.

McCartney, Scott, "Flier auctions: Better seats, going once, going twice...," *The Wall Street Journal*, 2013.

Menino, Thomas M., "State of the city address," https://www.cityofboston.gov/Images_Documents/State_of_the_City_2012_tcm3-30137.pdf 2012.

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," *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*, 2013, *103*, 80–106.

Peranson, Elliott, "Practical considerations in operating a match," 2005. Presented at the Jerusalem School in Economic Theory.

Roth, Alvin E. and Marilda Sotomayor, "The college admissions problem revisited," *Econometrica*, 1989, *57*, 559–570.

Service Academy Forums, "How does ROTC cadet receive designated branch?," 2012. `http://www.serviceacademyforums.com/showthread.php?t=26435`.

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

_ **and Tobias B. Switzer**, "Matching with (branch-of-choice) contracts at United States Military Academy," *Econometrica*, 2013, *81*, 451–488.

Westkamp, Alexander, "An analysis of the German university admissions system," *Economic Theory*, forthcoming.

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 $D_b^0 \equiv \varnothing$, and let $A_b^1 \equiv \varnothing$.

2. For each $\ell = 1, 2, \ldots$:

   (a) Let i be the $\sqsupset_\ell$-maximal agent $i \in I$ such that $i \notin \mathrm{i}(\cup_{b \in B} D_b^{\ell-1})$ 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 $b \equiv \mathrm{b}(x)$. 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 $z \in A_{\mathrm{b}(z)}^\ell$ but $z \notin D_{\mathrm{b}(z)}^{\ell-1}$ for some ℓ.

B Proofs Omitted from the Main Text

Proof of Proposition 1

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

Lemma B.1. *Suppose that $Y \subseteq Y' \subseteq X_b$, $|Y| = |\mathrm{i}(Y)|$, and $|Y'| = |\mathrm{i}(Y')|$. 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 $Y_{\mathrm{i}(x)} = \{x\}$ for each $x \in Y$ and that $Y'_{\mathrm{i}(x')} = \{x'\}$ 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 $x' \equiv \max_{\Pi^{s_b^\ell}} \left(V_b^\ell(Y')\right)$. If $x' \in V_b^\ell(Y)$, then clearly $x' = \max_{\Pi^{s_b^\ell}} \left(V_b^\ell(Y)\right)$; hence,

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

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

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

The claim implies that

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

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

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

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 weakly substitutable.

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 that blocks $\varpi(\tilde{Y})$. We fix some $b \in \mathrm{b}(Z)$, 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 $zP^{\mathrm{i}(z)}(\varpi(\tilde{Y}))_{\mathrm{i}(z)}$ 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 mehanism 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 $\mathrm{b}(x^\ell) = b = \mathrm{b}(x^{\ell'})$, 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

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

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

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

$$z\Pi^{s}(\varpi(\tilde{Y}))_{s} = Y_{s}. \tag{9}$$

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 $s' \in S_{\mathrm{b}(z)}$ with $s' \rhd^{\mathrm{b}(z)} s$ is assigned z in the computation of $C^{\mathrm{b}(z)}(A^{\ell+1}_{\mathrm{b}(z)}) = D^{\ell}_{\mathrm{b}(z)}$. (Such a step ℓ must exist since $z \notin Y$.) We let x^{ℓ} be the contract assigned to s in the computation of $C^{\mathrm{b}(z)}(A^{\ell+1}_{\mathrm{b}(z)})$. 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 $C^{\mathrm{b}(z)}(A^{\ell'+1}_{\mathrm{b}(z)})$ 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 $Y_{s}\Pi^{s}z$, contradicting (9). Thus, there cannot be a slot-contract $\langle z; s\rangle$ which slot-blocks $\tilde{Y}$. ◻

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 $s \in S_{\mathrm{b}(z)}$ 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 $C^{\mathrm{b}(z)}(A^{\ell+1}_{\mathrm{b}(z)})$.

Now, as $z \notin C^{\mathrm{b}(z)}(A^{\ell+1}_{\mathrm{b}(z)})$ and x is assigned to s in the computation of $C^{\mathrm{b}(z)}(A^{\ell+1}_{\mathrm{b}(z)})$, we know that $x\Pi^{s}z$. Moreover, as z is the first contract in $\varpi(\tilde{Z})$ to be rejected, we know that $xR^{\mathrm{i}(x)}(\varpi(\tilde{Z}))_{\mathrm{i}(x)}$. We say that because x has these two properties, x **supercedes** z **at** s.

If $x \neq (\varpi(\tilde{Z}))_{\mathrm{i}(x)}$, then we must have $xP^{\mathrm{i}(x)}(\varpi(\tilde{Z}))_{\mathrm{i}(x)}$ and it follows immediately that $\langle x; s\rangle$ slot-blocks $\tilde{Z}$. Meanwhile, if $x = (\varpi(\tilde{Z}))_{\mathrm{i}(x)}$, then there is some contract of the form

$\langle x; s' \rangle \in \tilde{Z}$, with $s' \neq s$. If $s \rhd^b s'$, then once again $\langle x; s \rangle$ slot-blocks $\tilde{Z}$, as $\langle x; s \rangle \, \tilde{\Pi}^s \, \langle z; s \rangle$ and $\langle x; s \rangle \, \tilde{P}^i \, \langle x; s' \rangle$ by construction. Finally, if $s' \rhd^b s$, then we let y be the contract assigned to s' in the computation of $C^{\mathrm{b}(x)}(A^{\ell+1}_{\mathrm{b}(x)})$. As $s' \rhd^b s$ and x is not assigned to s' in the computation of $C^{\mathrm{b}(x)}(A^{\ell+1}_{\mathrm{b}(x)})$, we must have $y \Pi^{s'} x$. Moreover, as z is the first contract in $\varpi(\tilde{Z})$ to be rejected, we know that $y R^{\mathrm{i}(y)} (\varpi(\tilde{Z}))_{\mathrm{i}(y)}$. Using the same terminology as before, we see that y supercedes x at s'.

Iterating these arguments, we see that either $\tilde{Z}$ must be slot-blocked, or there is an arbitrarily long sequence $s, s', \ldots$ of slots (in $S_{\mathrm{b}(z)}$) at which contracts are superceded. The former possibility cannot occur because $\tilde{Z}$ is slot-stable; the latter is impossible as well, because S is finite. $\square$

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-stable slot-outcome, by Theorem 3 of Hatfield and Milgrom (2005).

Proof of Theorem 3

The result is an immediate consequence of Theorem 2: Suppose that there exists a slot-stable outcome $\tilde{Y}$ that corresponds (under ϖ) to an agent-optimal stable outcome $Y = \varpi(\tilde{Y})$. By Theorem 2, the cumulative offer process produces an outcome Z that is stable. Moreover, for any slot-stable $\tilde{Z} \subseteq \tilde{X}$, each agent (weakly) prefers Z to $\varpi(\tilde{Z})$. Thus, we see that we must have $Z R^i \varpi(\tilde{Y}) = Y$ for each $i \in I$. If $Z \neq Y$, then, we have $Z P^i \varpi(\tilde{Y}) = Y$ for some $i \in I$, so that Z is stable and Pareto dominates Y—a contradiction.

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

$$\Phi_{\Pi}\left(\bar{P}^{I'}, P^{-I'}\right) P^{i} \Phi_{\Pi}\left(P^{I'}, P^{-I'}\right) \tag{10}$$

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

$$\varpi(\tilde{Z}) = \Phi_{\Pi}\left(\bar{P}^{I'}, P^{-I'}\right) P^{i} \Phi_{\Pi}\left(P^{I'}, P^{-I'}\right) = \varpi(\tilde{Y})$$

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).[55]

Proof of Theorem 5

To see this, we fix an agent i and let $\bar{\Pi}$ be an unambiguous improvement over Π 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,

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

that is, the order obtained from $\sqsupset$ by moving i to the bottom of each linear order $\sqsupset_{\ell}$.

[55]Theorem 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 substitutability condition and the law of aggregate demand hold automatically (for all preferences) in one-to-one matching markets like 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.[56] 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,

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

The substitutability of each branch choice function C^b in the presence of agent types then follows from Proposition 1, as weak substitutability 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 $\max_{\Pi^{s_b^\ell}} V_b^\ell(Y') = \emptyset_{s_b^\ell}$; hence,

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

[56] In the cumulative offer process associated to $\sqsupset'$, any contract x' with $\mathrm{i}(x') = i$ 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 $\bar{\Pi}$. Thus, such x' must be rejected under priorities Π, as otherwise $\bar{\Pi}$ would not be an unambiguous improvement over Π for i.

Proof of Proposition 4

We suppose that $Y \subseteq X$ is stable and consider the set of slot-contracts

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

Clearly, $\varpi(\tilde{Y}) = Y$. Thus, the desired result follows directly from the following claim.

Claim. *The slot-outcome $\tilde{Y}$ is slot-stable.*

Proof. We suppose to the contrary that $\tilde{Y}$ is slot-blocked, and let $\langle x; s \rangle$ be a slot-block of $\tilde{Y}$. We let y be the contract assigned to s in the computation of $C^{\mathrm{b}(x)}(Y)$, and let $z = Y_{\mathrm{i}(x)}$. As $\langle x; s \rangle$ slot-blocks $\tilde{Y}$, we have

$$\langle x; s \rangle \, \tilde{\Pi}^s \, \langle y; s \rangle \qquad \Longrightarrow \qquad x \Pi^s y. \tag{12}$$

If $z = \emptyset_{\mathrm{i}(x)}$ or $\mathrm{b}(z) \neq \mathrm{b}(x)$, the fact that $\langle x; s \rangle$ slot-blocks $\tilde{Y}$ implies that

$$\langle x; s \rangle \, \tilde{P}^{\mathrm{i}(x)} \, \langle z; s' \rangle \ \ (\text{for all } s' \in S_{\mathrm{b}(z)}) \qquad \Longrightarrow \qquad x P^{\mathrm{i}(x)} z;$$

hence $\{x\}$ blocks Y.

If $z \neq \emptyset_{\mathrm{i}(x)}$ and $\mathrm{b}(z) = \mathrm{b}(x)$, then we must have $z = x$ because $\mathrm{i}(z) = \mathrm{i}(x)$ and the market has agent types. Thus, there is some slot $s' \in S_{\mathrm{b}(x)}$ assigned $x = z$ in the computation of $C^{\mathrm{b}(x)}(Y)$. As $\langle x; s \rangle$ slot-blocks $\tilde{Y}$, we have

$$\langle x; s \rangle \, \tilde{P}^i \, \langle x; s' \rangle \, ;$$

hence $s \rhd^{\mathrm{b}(x)} s'$ by construction of $\tilde{P}^i$. But this, combined with (12), contradicts the construction of $C^{\mathrm{b}(x)}$, as it means that x should be assigned to s (or some higher-precedence slot) in the computation of $C^{\mathrm{b}(x)}(Y)$, rather than s'. $\square$

Proof of Proposition 6

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

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

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

We assume that

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

Claim. *The outcomes*

$$Y \equiv \{(i^m, b, \mathsf{t}_0) : m < \ell\} \cup \{(i^m, b, \mathsf{t}_+) : \ell \leq m < \ell'\} \cup \{(i^m, b, \mathsf{t}_0) : \ell' \leq m \leq |S_b|\},$$

$$Y' \equiv \{(i^m, b, \mathsf{t}_0) : m < \ell\} \cup \{(i^m, b, \mathsf{t}_+) : \ell < m \leq \ell'\} \cup \{(i^\ell, b, \mathsf{t}_0)\} \cup \{(i^m, b, \mathsf{t}_0) : \ell' < m \leq |S_b|\}$$

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 cadets i^m for whom $m < \ell$ or $m \geq \ell'$ hold their most-preferred contracts. Thus, any set blocking Y must be a subset of $Z^* \equiv \{(i^m, b, \mathsf{t}_0) : \ell \leq m < \ell'\}$. However, the contracts (i^m, b, t_+) are assigned to slots s_b^m ($\ell \leq m < \ell'$) in the computation of $C^b(Y)$, and $(i^m, b, \mathsf{t}_+) \Pi^{s_b^m} (i, b, \mathsf{t}_0)$ for all $i \in I$ and m with $\ell \leq m < \ell'$. 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 $B = \{b, b'\}$, $I = \{i, i', i'', j\}$, $\mathrm{i}(h_b) = h = \mathrm{i}(h_{b'})$ for each $h \in I$, and $\mathrm{b}(h_{b''}) = b''$ for each $h \in I$ and $b'' \in B$. We suppose that there are two types of agents—$\mathsf{T} = \{\mathsf{i}, \mathsf{j}\}$—and that $\mathrm{t}(i) = \mathrm{t}(i') = \mathrm{t}(i'') = \mathsf{i}$, while $\mathrm{t}(j) = \mathsf{j}$. We suppose that agents have preferences

$$
\begin{aligned}
P^{i} &: i_b \succ i_{b'} \succ \emptyset_i, \\
P^{i'} &: i'_b \succ \emptyset_{i'}, \\
P^{i''} &: i''_{b'} \succ \emptyset_{i''} \\
P^{j} &: j_b \succ j_{b'} \succ \emptyset_j.
\end{aligned}
$$

We suppose further that b has two slots, $s^1_b \rhd^b s^2_b$, with priorities given by

$$
\begin{aligned}
\Pi^{s^1_b} &: i_b \succ i'_b \succ i''_b \succ j_b \succ \emptyset_{s^1_b}, \\
\Pi^{s^2_b} &: i_b \succ j_b \succ i'_b \succ i''_b \succ \emptyset_{s^2_b},
\end{aligned}
$$

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

$$
\Pi^{s^1_{b'}} : i_b \succ j_b \succ i'_b \succ i''_b \succ \emptyset_{s^2_b}.
$$

In this example, the cumulative offer process outcome is $\{i_b, j_b, i''_{b'}\}$. If the precedence of slots s^1_b and s^2_b 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 $\mathsf{i} = \mathrm{t}(i'')$.

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 (forthcoming, 2012).

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 $x' \equiv \max_{\Pi^{s_b^\ell}} \left(V_b^\ell(Y \cup \{z, z'\})\right)$. As $z' \notin C^b(Y \cup \{z, z'\})$, we must have $x' \neq z'$; hence,

$$x' = \max_{\Pi^{s_b^\ell}} \left(V_b^\ell(Y \cup \{z, z'\}) \setminus \{z'\}\right) = \max_{\Pi^{s_b^\ell}} \left(V_b^\ell(Y \cup \{z\})\right) \tag{13}$$

by the inductive hypothesis. It then follows that

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

where the second equality follows from the inductive hypothesis, and the fourth equality follows from (13). This completes our induction. □

The claim implies the desired result, as it shows that

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

Now, we suppose that $\mathrm{i}(z), \mathrm{i}(z') \notin \mathrm{i}(Y)$, 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 $\mathrm{i}(z) \neq \mathrm{i}(z')$, as no branch ever selects two contracts with the same agent.

Claim. *For each ℓ with $1 \leq \ell \leq q_b$, let z^ℓ and y^ℓ be the contracts assigned to s_b^ℓ in the computation of $C^b(Y \cup \{z, z'\})$ and $C^b(Y \cup \{z\})$, respectively. We have $z^\ell \Gamma^{s_b^\ell} y^\ell$.*

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 $z^1 \Gamma^{s_b^\ell} y^1$ and $\mathrm{i}(H_b^1(Y \cup \{z, z'\})) \subseteq \mathrm{i}(H_b^1(Y \cup \{z\}) \cup \{z'\})$. Thus, we suppose that $z^{\ell'} \Gamma^{s_b^{\ell'}} y^{\ell'}$ and $\mathrm{i}(H_b^{\ell'}(Y \cup \{z, z'\})) \subseteq \mathrm{i}(H_b^{\ell'}(Y \cup \{z\}) \cup \{z'\})$ for all $\ell' < \ell$.

We have

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

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

$$
\begin{aligned}
\left((Y \cup \{z\}) \setminus \left(Y_{\mathrm{i}(H_b^{\ell-1}(Y \cup \{z\}))}\right)\right) &= \left((Y \cup \{z\}) \setminus \left(Y_{\mathrm{i}(H_b^{\ell-1}(Y \cup \{z\}) \cup \{z'\})}\right)\right) \\
&\subseteq \left((Y \cup \{z\}) \setminus \left(Y_{\mathrm{i}(H_b^{\ell-1}(Y \cup \{z, z'\}))}\right)\right) \\
&\subseteq \left((Y \cup \{z, z'\}) \setminus \left(Y_{\mathrm{i}(H_b^{\ell-1}(Y \cup \{z, z'\}))}\right)\right), \qquad (14)
\end{aligned}
$$

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

$\{z\}) \cup \{z'\})$.

The inclusion (14) implies that $z^\ell \Gamma^{s_b^\ell} y^\ell$. This observation completes the first part of the induction. Moreover, it quickly yields the second part. To see this, we observe that if $\mathrm{i}(z^\ell) \notin \mathrm{i}(H_b^{\ell-1}(Y \cup \{z\}))$, then either $z^\ell = z'$ or $z^\ell = y^\ell$, as $z^\ell \Gamma^{s_b^\ell} y^\ell$ and $\mathrm{i}(z') \notin \mathrm{i}(Y \cup \{z\})$. In either case, we have $\mathrm{i}(H_b^\ell(Y \cup \{z, z'\})) \subseteq \mathrm{i}(H_b^\ell(Y \cup \{z\}) \cup \{z'\})$. And finally, if $\mathrm{i}(z^\ell) \in \mathrm{i}(H_b^{\ell-1}(Y \cup \{z\}))$, then

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

so the induction is complete.[57] $\square$

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.

[57] Here, the first inclusion follows from the inductive hypothesis.