

Matching with Slot-Specific Priorities: Theory*

Scott Duke Kominers
Society of Fellows
Harvard University

Tayfun Sönmez
Department of Economics
Boston College

June 10, 2015

Abstract

We introduce a two-sided, many-to-one matching with contracts model in which agents with unit demand match to branches that 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, according to an order of precedence. 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 the slot-specific priorities framework and can be found by a cumulative offer mechanism that is strategy-proof and respects unambiguous improvements in priority.

JEL classification: C78, D47, D63, D78.

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

*The results in this paper were originally circulated as part of the earlier working paper “Designing for Diversity in Matching,” an extended abstract of which appeared in the *Proceedings of the 14th ACM Conference on Electronic Commerce*. The authors appreciate the helpful suggestions of Samson Alva, Nick Arnosti, 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, Parag Pathak, Philip Reny, Assaf Romm, Alvin Roth, Hugo Sonnenschein, Peter Troyan, Harald Uhlig, Utku Ünver, Alexander Westkamp, E. Glen Weyl, M. Bumin Yenmez, Hanzhe Zhang, members of the “Inequality: Measurement, Interpretation, and Policy” Working Group, numerous seminar audiences, two anonymous referees, and the editor, George Mailath. Kominers gratefully acknowledges the support of National Science Foundation grants CCF-1216095 and SES-1459912, an AMS–Simons travel grant, the Harvard Milton Fund, 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 and Microsoft Research New England. Sönmez gratefully acknowledges the support of Goldman Sachs Gives via Dalinc Ariburnu–Goldman Sachs Faculty Research Fund and National Science Foundation Grant SES-1426566.

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 *stable*, guaranteeing that no student ever envies a student with lower priority, and 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 (including 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, implementing affirmative action via matching market design typically causes agents’ priorities to vary across a given institution’s slots. A similar issue arises in the context of United States cadet–branch matching: there, some (but not all) slots at each branch of service grant increased priority for cadets willing to bid additional years of service. In this paper, we develop a general framework for handling these sorts of *slot-specific priority* structures.³ Our model 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.^{4,5}

¹These reforms include assignment of high school students in New York City (Abdulkadiroğlu et al. (2005b, 2009)), assignment of K–12 students to public schools in Boston (Abdulkadiroğlu et al. (2005a)), assignment of high school students to selective enrollment schools in Chicago (Pathak and Sönmez (2013)), and assignment of K–12 students to public schools in Denver. 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)).

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

³Like our model, the Westkamp (2013) model of *matching with complex constraints* permits priorities to vary across slots (see also Braun et al. (2014)). While Westkamp (2013) allows more general forms of interaction across slots than we allow in the present work, Westkamp (2013) 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)). In part motivated by our work, Echenique and Yenmez (forthcoming) have axiomatically characterized a class of substitutable priority/choice rules that allow schools to express preferences for diversity.

⁴The 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)).

⁵While our model embeds existing approaches to upper bound constraints in affirmative action (majority quotas, *a la* Kojima (2012), and minority reserves, *a la* Hafalir et al. (2013)), it does not incorporate *lower bound* quota constraints. Recent work of Ehlers et al. (2014) and Fragiadakis et al. (2014) has shown how to integrate lower bound constraints into matching.

We show how markets with slot-specific priorities can be cleared by 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 the cumulative offer mechanism is strategy-proof.⁶ In markets with slot-specific priorities, agent-optimal stable outcomes may not exist; nevertheless, as we show, the cumulative offer mechanism remains strategy-proof. 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.⁷

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.

Our paper also has a methodological contribution: In general, slot-specific priorities fail the substitutability condition that has so far been key in the 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 (2006); Hatfield and Kominers (2014)). 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)).⁸ 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;⁹ moreover, we show that the cumulative offer mechanism in the “true” matching

⁶Here, by *agent-optimal stable outcomes*, we mean stable outcomes that all agents prefer to all other stable outcomes.

⁷The stability conclusion can be derived by combining the fact that slot-specific priorities induce bilaterally substitutable choice functions (Lemma 1) with prior results of Hatfield and Kojima (2010). Our other results do not follow from prior work; they depend upon structure present in our specific model that also leads to a direct proof of the stability result (see the discussion at the end of Section 2.2.1).

⁸Thus, in particular, our model falls outside of the domains which Echenique (2012) and Schlegel (2015) have shown can be handled with only the Kelso and Crawford (1982) *matching with salaries* framework (see also Kominers (2012)).

⁹The converse result is not true, in general—there may be stable outcomes which are not associated to slot-stable outcomes.

market gives the outcome that 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 acquisition 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 stable, strategy-proof, and (unambiguous-)improvement-respecting.

The remainder of this paper is organized as follows. We present our model of matching with slot-specific priorities in Section 2. Then, in Section 3, we introduce the agent–slot matching market and derive key properties of the cumulative offer mechanism. In Section 4, we present an application to airline seat upgrade allocation. Section 5 concludes. All proofs are presented in the Appendix.

2 Model

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 . Here, a *branch* could represent, for example, a branch of the military (as in cadet–branch matching), or a school (as in school choice). Each contract $x \in X$ is between an agent $i(x) \in I$ and branch $b(x) \in B$.¹⁰ We extend the notations $i(\cdot)$ and $b(\cdot)$ to sets of contracts $Y \subseteq X$ by setting $i(Y) \equiv \cup_{y \in Y} \{i(y)\}$ and $b(Y) \equiv \cup_{y \in Y} \{b(y)\}$. For $Y \subseteq X$, we denote $Y_i \equiv \{y \in Y : i(y) = i\}$ and $Y_{I'} \equiv \cup_{i \in I'} Y_i$; analogously, we denote $Y_b \equiv \{y \in Y : b(y) = b\}$ and $Y_{B'} \equiv \cup_{b \in B'} Y_b$.

Each agent $i \in I$ has a (linear) preference order P^i (with weak order R^i) over contracts in $X_i = \{x \in X : i(x) = i\}$. 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 ; we use the convention that $\emptyset_i P^i x$ if $x \in X \setminus X_i$. We say that the contracts $x \in X$ for which $\emptyset_i P^i x$ are **unacceptable to i** . We denote the profile of all agents’ preferences by P .

Each branch $b \in B$ has a set S_b of slots; each slot can be assigned at most one contract in

¹⁰A 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.

$X_b \equiv \{x \in X : b(x) = b\}$. 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.¹¹ We set $S \equiv \cup_{b \in B} S_b$ and denote the profile of all slots’ priorities by Π .

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

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 $y \in Y_s$.¹²

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 best 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 $|S_b|$ 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** \triangleright^b . We denote $S_b \equiv \{s_b^1, \dots, s_b^{q_b}\}$ with $q_b \equiv |S_b|$ and the understanding that $s_b^\ell \triangleright^b s_b^{\ell+1}$ unless otherwise noted. The interpretation of \triangleright^b is that if $s \triangleright^b s'$ then—whenever possible—branch b fills slot s before filling s' . Formally, the choice $C^b(Y)$ of a branch $b \in B$ from contract set $Y \subseteq X$ is defined as follows:

- First, slot s_b^1 is assigned the contract x^1 that is $\Pi^{s_b^1}$ -maximal among contracts in Y .
- Then, slot s_b^2 is assigned the contract x^2 that is $\Pi^{s_b^2}$ -maximal among contracts in the set $Y \setminus Y_{i(x^1)}$ of contracts in Y with agents other than $i(x^1)$.
- This process continues in sequence, with each slot s_b^ℓ being assigned the contract x^ℓ that is $\Pi^{s_b^\ell}$ -maximal among contracts in the set $Y \setminus Y_{i(\{x^1, \dots, x^{\ell-1}\})}$.¹³

2.1 Solution Concept

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

- Y contains at most one contract for each agent, i.e. $|Y_i| \leq 1$ for each $i \in I$, and

¹¹As with agents, we use the convention that $\emptyset_s \Pi^s x$ if $x \in X \setminus X_s$.

¹²Here, 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 .

¹³If 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 $\emptyset_{s_b^\ell}$.

- Y contains at most q_b contracts for each branch b , i.e. $|Y_i| \leq q_b$ for each $b \in B$.

We follow the Gale and Shapley (1962) tradition in focusing on outcomes that are *stable* in the sense that (1) neither agents nor branches wish to unilaterally walk away from their assignments, and (2) agents and branches cannot benefit by recontracting outside of the assigned outcome. Formally, we say that an outcome Y is **stable** if it is

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

Individual rationality for agents requires that no agent be assigned a contract that he finds unacceptable. In our context, as in other matching with contracts models, individual rationality for branches corresponds to a form of “respect for branch choices.”¹⁴ Requiring unblockedness is a form of eliminating justified envy; it means that no agent desires a slot at which he has a justified claim—with some desirable contract—under the priority and precedence structure.¹⁵

2.2 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;¹⁶ thus, they help contextualize our results and illustrate some of the technical difficulties that arise in our general framework.

2.2.1 Substitutability Conditions

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

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

¹⁴Notably, in our context, individual rationality does not have the more classical connotation of “preventing matches to unacceptable partners,” because respecting the slot-specific priority structure (along with the given precedence structure) may require holding open a slot s even when contracts that are acceptable at some other slot $s' \neq s$ remain available. If we additionally require that all contracts that are acceptable at some slot of b be acceptable at all slots of b , then our individual rationality condition has the more classical meaning.

¹⁵We thank a referee for pointing out these distinctions.

¹⁶As 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.

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.¹⁷

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 (2015)) and the Ostrovsky (2008) supply chain matching framework (Hatfield and Kominers (2012)). However, substitutability is *not* necessary for the guaranteed existence of stable outcomes in settings where agents have unit demand (Hatfield and Kojima (2008, 2010)). Indeed, as Hatfield and Kojima (2010) showed, the following condition weaker than substitutability suffices not only for the existence of stable outcomes, but also to guarantee that there is no conflict of interest among agents.¹⁸

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

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

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

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 1. Let $X = \{x_1, x_2, y_2\}$, with $B = \{b\}$, $I = \{i, j\}$, $i(x_1) = i = i(x_2)$, $i(y_2) = j$, and $b(x_1) = b(x_2) = b(y_2) = b$. If b has two slots, $s_b^1 \triangleright^b s_b^2$, with priorities given by

$$\begin{aligned} \Pi^{s_b^1} &: x_1 \succ \emptyset_{s_b^1}, \\ \Pi^{s_b^2} &: x_2 \succ y_2 \succ \emptyset_{s_b^2}, \end{aligned}$$

¹⁷The 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 (2013)). This condition is naturally satisfied in most economic environments—including ours (see Lemma A.1 of Appendix A).

¹⁸As in the work of Hatfield and Milgrom (2005), an irrelevance of rejected contracts condition (which is naturally satisfied in our setting—see Footnote 17) is implicitly assumed throughout the work of Hatfield and Kojima (2010) (Aygün and Sönmez (2012)).

then C^b fails the unilateral substitutability condition: if we take $z = y_2$, $z' = x_1$, and $Y = \{x_2\}$, then $z = y_2 \notin C^b(\{x_2, y_2\}) = C^b(Y \cup \{z\})$, but $z = y_2 \in C^b(\{x_1, x_2, y_2\}) = C^b(Y \cup \{z, z'\})$, even though $i(z) = i(y_2) = j \notin \{i\} = i(\{x_2\}) = i(Y)$.

The choice functions C^b do, however, satisfy the bilateral substitutability condition introduced by Hatfield and Kojima (2010).

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

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

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

Lemma 1. *Every choice function C^b is bilaterally substitutable.*

Combining Lemma 1 with Theorem 1 of Hatfield and Kojima (2010) implies the existence of stable outcomes in our setting.¹⁹ However, the bilateral substitutability condition is not sufficient for the other key results necessary for matching market design (e.g., the existence of strategy-proof matching mechanisms); to obtain these additional results in our framework, we draw upon structure present in our specific model (see Section 3.1).²⁰

2.2.2 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).²¹

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

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

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 2. Let $X = \{x_1, x_2, y_1\}$, with $B = \{b\}$, $I = \{i, j\}$, $i(x_1) = i = i(x_2)$, $i(y_1) = j$, and $b(x_1) = b(x_2) = b(y_1) = b$. If b has two slots, $s_b^1 \triangleright^b s_b^2$, with priorities given by

$$\begin{aligned} \Pi^{s_b^1} &: x_1 \succ y_1 \succ \emptyset_{s_b^1}, \\ \Pi^{s_b^2} &: x_2 \succ \emptyset_{s_b^2}, \end{aligned}$$

¹⁹As discussed in Footnote 18, this logic implicitly requires an irrelevance of rejected contracts condition (Aygün and Sönmez (2012)).

²⁰The structure we identify also gives rise to a self-contained existence proof, which does not exploit the bilateral substitutability condition.

²¹Alkan (2002) and Alkan and Gale (2003) introduced a related *cardinal monotonicity* condition.

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

$$|C^b(\{x_2, y_1\})| = |\{x_2, y_1\}| = 2 > 1 = |\{x_1\}| = |C^b(\{x_1, x_2, y_1\})|.$$

3 Results

We now develop our main results: In Section 3.1, we associate our original market to a (one-to-one) matching market in which slots, rather than branches, compete for contracts. Next, in Section 3.2, 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 3.3, 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 3.4, we show that the mechanism that selects the cumulative offer process outcome is stable, strategy-proof, and improvement-respecting.

3.1 Associated Agent–Slot Matching Market

We now associate a (one-to-one) **agent–slot matching market** to our original market, by extending the contract set X to the set \tilde{X} defined by

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

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

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

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

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

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

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

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

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

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

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

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

Our next result shows that slot-stable slot-outcomes project to stable outcomes.

Lemma 2. *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 2 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 the **agent-optimal slot-stable mechanism**, i.e. the mechanism that selects the agent-optimal slot-stable slot-outcome in the agent–slot market.²²

3.2 The Cumulative Offer Process

We now introduce the cumulative offer process for matching with contracts, which generalizes the *agent-proposing deferred acceptance algorithm* of Gale and Shapley (1962).

²²Classic results for one-to-one matching with contracts show that this mechanism exists and is well-defined (Crawford and Knoer (1981); Kelso and Crawford (1982); Hatfield and Milgrom (2005)).

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

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

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

If at any time no agent is able to propose a new contract—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 the algorithm terminates.

In the cumulative offer process, agents propose contracts sequentially. Branches accumulate offers, at each step choosing 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 (Theorem A.1 of Appendix A). At the time that we first wrote our paper, there was no general order-independence result covering the slot-specific priority framework, although an order-independence result was known for settings with unilaterally substitutable branch choice functions (Hatfield and Kojima (2010)).²³ Since we first circulated our paper, however, Hirata and Kasuya (2014) have proven an order-independence result more general than ours: the cumulative offer process is order-independent whenever all branches have bilaterally substitutable choice functions that satisfy the irrelevance of rejected contracts condition.

We now show that the cumulative offer process outcome has a natural interpretation: it corresponds to the agent-optimal slot-stable slot-outcome in the agent–slot matching market.

²³As our Example 1 illustrates, slot-specific priorities may not induce unilaterally substitutable branch choice functions.

Lemma 3. *The agent-optimal slot-stable slot-outcome corresponds (under projection ϖ) to the outcome of the cumulative offer process.*

The proof of Lemma 3 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.²⁴ This observation implies that no contract held by a slot $s \in S$ at some step of the cumulative offer process has higher priority than the contract s holds at the end of the process; it follows that the cumulative offer process outcome Y is the ϖ -projection of a slot-stable slot-outcome \tilde{Y} . Then, we demonstrate that agents (weakly) prefer \tilde{Y} to the agent-optimal slot-stable slot-outcome \tilde{Z} , which exists by Theorem 3 of Hatfield and Milgrom (2005). It follows that $\tilde{Y} = \tilde{Z}$; as $\varpi(\tilde{Y}) = Y$, this proves Lemma 3.

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.²⁵ 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 Lemma 3, 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 2, that correspondence is far more delicate than the one obtained by Roth and Sotomayor (1989): it is not onto (see Example 3), and it depends crucially on the explicit incorporation of slot-precedence into agents’ extended preferences.

The cumulative offer process always terminates, as the set X is finite and the full set of contracts available, $\cup_{b \in B} A_b^\ell$, grows monotonically in ℓ . Moreover, it shows that the cumulative offer process outcome is independent of proposal order, stable, and somewhat distinguished among stable outcomes.

Theorem 1. *The cumulative offer process produces an outcome that 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 1 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

²⁴Here, 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 $C^b(A_b^{\ell+1})$.

²⁵In 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).

prefers the cumulative offer process outcome to *all* other stable outcomes; we demonstrate this explicitly in the next section.

3.3 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 $Y \subseteq X$ that Pareto dominates all other stable outcomes is called an **agent-optimal stable outcome**.²⁶ For general slot-specific priorities, agent-optimal stable outcomes need not exist, as the following example shows.

Example 3. Let $X = \{x_0, x_1, y_0, y_1, z_0, z_1\}$, with $B = \{b\}$, $I = \{i, j, k\}$, $i(x_0) = i = i(x_1)$, $i(y_0) = j = i(y_1)$, $i(z_0) = k = i(z_1)$ and $b(w) = b$ for each $w \in X$. We suppose that $x_0 P^i x_1 P^i \emptyset_i$, $y_0 P^j y_1 P^j \emptyset_j$, and $z_0 P^k z_1 P^k \emptyset_k$, and that b has two slots, $s_b^1 \triangleright^b s_b^2$, with slot priorities given by

$$\begin{aligned} \Pi^{s_b^1} &: x_1 \succ y_1 \succ z_1 \succ x_0 \succ y_0 \succ z_0 \succ \emptyset_{s_b^1}, \\ \Pi^{s_b^2} &: x_0 \succ x_1 \succ y_0 \succ y_1 \succ z_0 \succ z_1 \succ \emptyset_{s_b^2}. \end{aligned}$$

In this setting, the outcomes $Y \equiv \{y_1, x_0\}$ and $Y' \equiv \{x_1, y_0\}$ 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 1, the cumulative offer process yields 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 3.²⁷

Example 4. Let $X = \{x_0, x_1, x_*, y_0, y_1, y_*, z_0, z_1\}$, with $B = \{b\}$, $I = \{i, j, k\}$, $i(x_0) = i(x_1) = i(x_*) = i$, $i(y_0) = i(y_1) = i(y_*) = j$, and $i(z_0) = k = i(z_1)$.²⁸ We suppose that $i_0 P^i i_* P^i i_1 P^i \emptyset_h$, $y_0 P^j y_* P^j y_1 P^j \emptyset_h$, and $z_0 P^k z_1 P^k \emptyset_k$, and that b has two slots, $s_b^1 \triangleright^b s_b^2$, with slot priorities given by

$$\begin{aligned} \Pi^{s_b^1} &: y_1 \succ z_1 \succ x_0 \succ y_0 \succ z_0 \succ \emptyset_{s_b^1}, \\ \Pi^{s_b^2} &: x_* \succ y_* \succ x_0 \succ x_1 \succ y_0 \succ y_1 \succ z_0 \succ z_1 \succ \emptyset_{s_b^2}. \end{aligned}$$

In this setting, the outcome $Y \equiv \{y_1, x_0\}$ is the agent-optimal stable outcome. However, the outcome of the cumulative offer process is $Y'' \equiv \{y_1, x_*\}$, which is Pareto dominated by Y .

²⁶That 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 $Y' \neq Y$.

²⁷We thank Fuhito Kojima for this example.

²⁸Clearly (in order for $b(\cdot)$ to be well-defined), we must have $b(x) = b$ for each $x \in X$, as $|B| = 1$.

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 1.

Theorem 2. *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.*

3.4 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** $\varphi_{\Pi} : \times_{i \in I} \mathcal{S}^i \rightarrow X$ that selects an outcome for each choice of agent strategies. We confine our attention to **direct mechanisms**, i.e. mechanisms for which the strategy spaces correspond to the preference domains: $\mathcal{S}^i = \mathcal{P}^i$, where \mathcal{P}^i denotes the set of preference relations for agent $i \in I$. Direct 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 $\Phi_{\Pi} : \times_{i \in I} \mathcal{P}^i \rightarrow X$.

3.4.1 Stability and Strategy-Proofness

A mechanism φ is **stable** if it always selects an outcome stable with respect to slot priorities and input preferences. 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 $P^I \in \times_{j \in I} \mathcal{P}^j$, and $\bar{P}^i \neq P^i$ such that $\varphi(\bar{P}^i, P^{-i}) P^i \varphi(P^i, P^{-i})$. Similarly, a mechanism φ is **group strategy-proof** if there is no set of agents $I' \subset I$, preference profile $P^I \in \times_{j \in I} \mathcal{P}^j$, and $\bar{P}^{I'} \neq P^{I'}$ such that $\varphi(\bar{P}^{I'}, P^{-I'}) P^i \varphi(P^{I'}, P^{-I'})$ for all $i \in I'$.

It follows immediately from Theorem 1 that the cumulative offer mechanism is stable. Meanwhile, Theorem 1 of Hatfield and Kojima (2009) implies that the agent-optimal slot-stable mechanism is (group) strategy-proof in the agent-slot matching market. Thus, we see that the cumulative offer mechanism is (group) strategy-proof, as any $\bar{P}^{I'} \neq P^{I'}$ such that $\varphi(\bar{P}^{I'}, P^{-I'}) P^i \varphi(P^{I'}, P^{-I'})$ for all $i \in I'$ would give rise to a profitable manipulation

($\tilde{P}^{I'} \neq \tilde{P}^{I'}$) of the agent-optimal slot-stable mechanism. These observations are summarized in the following theorem.

Theorem 3. *The cumulative offer mechanism Φ_{Π} is stable and (group) strategy-proof.*

Theorem 3 generalizes two main results (Propositions 2 and 3) of Sönmez and Switzer (2013) (and the analogous results of Sönmez (2013)) for cadet–branch matching. The proof in the cadet–branch matching setting follows directly from results of Hatfield and Kojima (2010) and the fact that the priority structure in cadet–branch matching introduces unilaterally substitutable choice functions (Lemma 2 of Sönmez and Switzer (2013)). As the choice functions in our framework fail the unilateral substitutability condition, in general, we cannot rely upon the same approach here; this is why we develop our more direct approach using the agent–slot matching market. Theorem 3 also generalizes analogous findings of Hafalir et al. (2013) for the setting of affirmative action with minority reserves, as well as existence and strategy-proofness results from more classical matching with contracts frameworks (e.g., Hatfield and Milgrom (2005)).

3.4.2 Respect for Unambiguous Improvements

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

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.²⁹ Nevertheless, respect for improvements is essential in settings like

²⁹Respect for improvements is, however, of importance in the growing normative literature on school choice design. For example, Hatfield et al. (2015) have used this condition in analyzing how school choice mechanism selection can impact schools' incentives for self-improvement.

the airline seat upgrade application we introduce in Section 4—there, it implies that customers never want to decrease their frequent-flyer status levels.³⁰ Similarly, respect for unambiguous improvements is important in cadet–branch matching, where cadets can influence their priority rankings directly—and may (in the absence of respect for improvements) take perverse steps to lower their priorities.³¹

Theorem 4. *The cumulative offer mechanism Φ_{Π} respects unambiguous improvements.*

Theorem 4 generalizes Proposition 4 of Sönmez and Switzer (2013). To prove Theorem 4, we again work in the agent–slot matching market. Specifically, we use an argument closely analogous to that of Balinski and Sönmez (1999) to show that the agent-optimal slot-stable mechanism satisfies a condition analogous to respect of unambiguous improvements; Theorem 4 then follows from Lemma 3.³²

4 Application: 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.

³⁰Formally, for this claim we need the 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.

³¹As 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[...].?”

³²A 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 $x, x' \in (X_i \cap X_b)$,

$$\text{if } x \bar{\Pi}^s x' \text{ and } x' \Pi^s x, \text{ then } x P^i x'. \quad (1)$$

As Theorem 4 shows that Φ_{Π} respects unambiguous improvements, we see *a fortiori* that Φ_{Π} respects unambiguous improvements that satisfy the additional condition (1).

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.³³

Example 5. There is a unique branch b , the *business class*, which has two upgrade slots

³³Note that this example is closely analogous to Example 1.

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 .³⁴

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 $s_b^1 \triangleright^b s_b^2$, and

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

Observe that the resulting business class choice function C^b fails the unilateral substitutability condition (introduced on page 6): 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 $i(z) = i(j_m) = j \notin \{i\} = i(i_m) = i(Y)$.

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 recently begun implementing (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 where there are multiple media of exchange and slot-specific-priorities.³⁵

³⁴This way, if there is one cash-customer and one miles-customer, the cash-customer will not block the airline from assigning both upgrade seats.

³⁵One 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), as well as the work of Hatfield

5 Conclusion

We have introduced a model of many-to-one *matching with slot-specific priorities*. In our framework, branches have a precedence-ordered set of slots, and each slot has its own linear priority order over contracts. Although branches’ choice functions may not satisfy the substitutability conditions used throughout most of the work on matching with contracts, we find that in our setting the standard cumulative offer process finds stable outcomes, is strategy-proof, and respects unambiguous improvements in priority.

Our work shows that the existence of agent-optimal stable outcomes is not necessary for strategy-proof stable matching, and reinforces and expands the applicability of the matching with contracts framework. Additionally, our general model allows us to address novel market design applications like the allocation of airline seat upgrades.

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.³⁶ 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 Kominers (2014)). 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.

³⁶Unfortunately, we do not have a complete characterization of the set of preferences that have slot-specific priority structure. One priority structure of which we are aware that is not covered by our model is that used in the German university admission system (Westkamp (2013); Braun et al. (2014)). Our model also cannot embed lower-bound constraints of the type studied by Ehlers et al. (2014) and Fragiadakis et al. (2014).

- , – , **and** – , “The New York City high school match,” *American Economic Review*, 2005, *95*, 364–367.
- , – , **and** – , “Strategyproofness versus efficiency in matching with indifference: 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.
- 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,” 2013, *103* (5), 2050–2051. *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,” *Games and Economic Behavior*, 2014, *85*, 232–251.
- Crawford, Vincent P. and Elsie Marie Knoer**, “Job matching with heterogeneous firms and workers,” *Econometrica*, 1981, *49*, 437–450.
- Echenique, Federico**, “Contracts vs. salaries in matching,” *American Economic Review*, 2012, *102*, 594–601.
- **and Bumin Yenmez**, “How to control controlled school choice,” *American Economic Review*, forthcoming.
- **and Jorge Oviedo**, “A theory of stability in many-to-many matching markets,” *Theoretical Economics*, 2006, *1*, 233–273.

- Ehlers, Lars, Isa E. Hafalir, M. Bumin Yenmez, and Muhammed A. Yildirim**, “School choice with controlled choice constraints: Hard bounds versus soft bounds,” *Journal of Economic Theory*, 2014, *153*, 648–683.
- Fleiner, Tamás**, “A fixed point approach to stable matchings and some applications,” *Mathematics of Operations Research*, 2003, *28*, 103–126.
- Fragiadakis, Daniel, Atsushi Iwasaki, Peter Troyan, Suguru Ueda, and Makoto Yokoo**, “Strategy-proof mechanisms for two-sided matching with minimum and maximum quotas,” 2014. Mimeo, Stanford University.
- 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**, “Matching in networks with bilateral contracts,” *American Economic Journal: Microeconomics*, 2012, *4*, 176–208.
- **and** – , “Hidden substitutes,” 2014. Mimeo, Harvard Business School.
- **and** – , “Contract design and stability in many-to-many matching,” 2015. Harvard Business School Working Paper.
- , **Fuhito Kojima, and Yusuke Narita**, “Promoting school competition through school choice: A market design approach,” 2015. Stanford Graduate School of Business Working Paper.
- Hirata, Daisuke and Yusuke Kasuya**, “Cumulative offer process is order-independent,” *Economics Letters*, 2014, *124* (1), 37–40.

- 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.
- 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 Summer School in Economic Theory.
- Roth, Alvin E. and Marilda Sotomayor**, “The college admissions problem revisited,” *Econometrica*, 1989, *57*, 559–570.
- Schlegel, Jan Christoph**, “Contracts versus salaries in matching: A general result,” 2015.
- 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*, 2013, 53 (3), 561–589.

A Proofs Omitted from the Main Text

Proof of Lemma 1

We first prove a lemma which shows that branch choice functions satisfy the *irrelevance of rejected contracts* property of Aygün and Sönmez (2013, 2012).

Lemma A.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 of contracts not assigned to any slot of b by the end of 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 $x' \equiv \max_{\Pi^s_b} (V_b^\ell(Y \cup \{z, z'\}))$. As $z' \notin C^b(Y \cup \{z, z'\})$, we must have $x' \neq z'$; hence,

$$x' = \max_{\Pi^s_b} (V_b^\ell(Y \cup \{z, z'\}) \setminus \{z'\}) = \max_{\Pi^s_b} (V_b^\ell(Y \cup \{z\})) \quad (2)$$

by the inductive hypothesis. It then follows that

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

where the second equality follows from the inductive hypothesis, and the fourth equality follows from (2). This completes our induction. \square

The claim implies the desired result, as it shows that

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

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

Claim. 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 of contracts assigned to slots of b at the end of step ℓ of the computation of $C^b(Z)$. We proceed by double induction: Clearly, either $z^1 = \max_{\Pi^{s_b^1}}(Y \cup \{z, z'\}) = \max_{\Pi^{s_b^1}}(Y \cup \{z\}) = y^1$ or $z^1 = \max_{\Pi^{s_b^1}}(Y \cup \{z, z'\}) = z'$, so $z^1 \Gamma^{s_b^1} y^1$ and $i(H_b^1(Y \cup \{z, z'\})) \subseteq i(H_b^1(Y \cup \{z\}) \cup \{z'\})$. Thus, we suppose that $z^{\ell'} \Gamma^{s_b^{\ell'}} y^{\ell'}$ and $i(H_b^{\ell'}(Y \cup \{z, z'\})) \subseteq i(H_b^{\ell'}(Y \cup \{z\}) \cup \{z'\})$ for all $\ell' < \ell$.

We have

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

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

$$\begin{aligned} \left(\left((Y \cup \{z\}) \setminus \left(Y_{i(H_b^{\ell-1}(Y \cup \{z\}))}\right)\right)\right) &= \left(\left((Y \cup \{z\}) \setminus \left(Y_{i(H_b^{\ell-1}(Y \cup \{z\}) \cup \{z'\})}\right)\right)\right) \\ &\subseteq \left(\left((Y \cup \{z\}) \setminus \left(Y_{i(H_b^{\ell-1}(Y \cup \{z, z'\}))}\right)\right)\right) \\ &\subseteq \left(\left((Y \cup \{z, z'\}) \setminus \left(Y_{i(H_b^{\ell-1}(Y \cup \{z, z'\}))}\right)\right)\right), \end{aligned} \quad (3)$$

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

The inclusion (3) 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 $i(z^\ell) \notin i(H_b^{\ell-1}(Y \cup \{z\}))$, then either $z^\ell = z'$ or $z^\ell = y^\ell$, as $z^\ell \Gamma^{s_b^\ell} y^\ell$ and $i(z') \notin i(Y \cup \{z\})$. In either case, we have $i(H_b^\ell(Y \cup \{z, z'\})) \subseteq i(H_b^\ell(Y \cup \{z\}) \cup \{z'\})$. And finally, if $i(z^\ell) \in i(H_b^{\ell-1}(Y \cup \{z\}))$, then

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

so the induction is complete.³⁷ □

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 $C^b(Y \cup \{z\})$ would have to 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.

Proof of Lemma 2

Claim. *The outcome $\varpi(\tilde{Y})$ is individually rational.*

Proof. If there were some $i \in I$ such that $C^i(\varpi(\tilde{Y})) \neq (\varpi(\tilde{Y}))_i$, then $(\varpi(\tilde{Y}))_i$ must be unacceptable to i ; hence, all slot-contracts of the form $\langle (\varpi(\tilde{Y}))_i; s \rangle$ must be unacceptable to i under \tilde{P}^i —impossible, as some such contract is in \tilde{Y} , and \tilde{Y} is individually rational for agents and slots.

Now, we suppose that there were some $b \in B$ such that $C^b(\varpi(\tilde{Y})) \neq (\varpi(\tilde{Y}))_b$. In this case, there is at least one contract $y \in (\varpi(\tilde{Y}))_b$ such that $\langle y; s \rangle \in \tilde{Y}$ for some slot $s \in S_b$, but y is not assigned to s in the computation of $C^b(\varpi(\tilde{Y}))$.³⁸ We let $\langle y; s \rangle$ be the slot-contract such that y is not assigned to s in the computation of $C^b(\varpi(\tilde{Y}))$ for which s is ranked highest under the precedence order \triangleright^b .

Now, as \tilde{Y} is individually rational for agents and slots, we know that $\langle y; s \rangle \tilde{\Pi}^s \emptyset_s$; hence, we must have $y\Pi^s \emptyset_s$. Then, as y is not assigned to s in the computation of $C^b(\varpi(\tilde{Y}))$ and s is the highest-precedence slot whose assignment in the computation of $C^b(\varpi(\tilde{Y}))$ does not correspond to its slot-contract in \tilde{Y} , we see that there must be some contract $z \neq y$ that is assigned to s in the computation of $C^b(\varpi(\tilde{Y}))$. But now, we must have $z\Pi^s y$, which implies that $\langle z; s \rangle \tilde{\Pi}^s \langle y; s \rangle$. As $z \in \varpi(\tilde{Y})$, there is some slot-contract of the form $\langle z; s' \rangle \in \tilde{Y}$. As z is available to be assigned to s in the computation of $C^b(\varpi(\tilde{Y}))$, we must have $s \triangleright^b s'$ —as otherwise, we would have $s' \triangleright^b s$ and z is not assigned to s' in the computation of $C^b(\varpi(\tilde{Y}))$, contradicting our precedence-maximality assumption on s . But then, since $s \triangleright^b s'$, by construction, we must have $\langle z; s \rangle \tilde{P}^{i(z)} \langle z; s' \rangle$. As we also have $\langle z; s \rangle \tilde{\Pi}^s \langle y; s \rangle$, we see that $\langle z; s \rangle$ is a slot-block of \tilde{Y} —contradicting the fact that \tilde{Y} is slot-stable. □

³⁷Here, the first inclusion follows from the inductive hypothesis.

³⁸In particular, as $C^b(\varpi(\tilde{Y})) \neq (\varpi(\tilde{Y}))_b$, the set of contracts

$$\left\{ \langle x; s \rangle \in \tilde{Y} : x \in [(\varpi(\tilde{Y}))_b \setminus C^b(\varpi(\tilde{Y}))] \right\}$$

is nonempty; this implies the existence of the claimed $\langle y; s \rangle$.

Claim. *The outcome $\varpi(\tilde{Y})$ is unblocked.*

Proof. For $\tilde{Y} \subseteq \tilde{X}$, suppose that $Z \subseteq X$ is a set of contracts that blocks $\varpi(\tilde{Y})$. We fix some $b \in \mathfrak{b}(Z)$, and observe that there must be a contract $z \in Z_b \setminus \varpi(\tilde{Y})$ that is assigned to the highest-precedence slot, s_b^ℓ , among those slots that 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 that is assigned to slot s_b^ℓ in the computation of $C^b(\varpi(\tilde{Y}))$.

It is clear that $z \Pi^{s_b^\ell} x$, by construction. Thus, we have $\langle z; s_b^\ell \rangle \tilde{\Pi}^{s_b^\ell} \langle x; s_b^\ell \rangle$. Meanwhile, we know that $z P^{i(z)}(\varpi(\tilde{Y}))_{i(z)}$ because Z blocks $\varpi(\tilde{Y})$. It follows that $\langle z; s_b^\ell \rangle$ is a slot-block for \tilde{Y} , contradicting the fact that \tilde{Y} is slot-stable. \square

Proof of Lemma 3

We denote by $D_b^{\ell-1}$ and A_b^ℓ 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 $z \in A_{\mathfrak{b}(z)}^\ell$ but $z \notin D_{\mathfrak{b}(z)}^\ell$ for some ℓ .

We consider the cumulative offer process in terms of a specific **proposal order**

$$\sqsupset = (\sqsupset_\ell)_{\ell=1}^\infty,$$

which specifies an ordering of agents \sqsupset_ℓ for each step of the cumulative offer process. The agent who proposes in step ℓ of the cumulative offer process under proposal order \sqsupset is the \sqsupset_ℓ -maximal agent who both (1) does not have a contract on hold at any branch and (2) wants to propose a new contract.³⁹

We prove the following result, which is slightly more general than Lemma 3.

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

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

We now prove two simple lemmatta that together show that slots' assigned contracts improve (with respect to slot priorities) over the course of the cumulative offer process.

Lemma A.2. *Suppose that $Y \subseteq Y' \subseteq X_b$, $|Y| = |i(Y)|$, and $|Y'| = |i(Y')|$. Then, if $y \in Y$ and $y' \in Y'$ are the contracts assigned to $s \in S_b$ in the computations of $C^b(Y)$ and $C^b(Y')$, respectively, we have $y' \Gamma^s y$.*

³⁹If no such agent exists, then the process terminates.

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

Claim. *Let $V_b^\ell(Z)$ denote the set of contracts not assigned to any slot of b by the end of step $\ell - 1$ of the computation of $C^b(Z)$. Then, $V_b^\ell(Y) \subseteq V_b^\ell(Y')$.*

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

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

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

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

The claim implies that

$$\left(\max_{\Pi^{s_b^\ell}} V_b^\ell(Y') \right) \Gamma^{s_b^\ell} \left(\max_{\Pi^{s_b^\ell}} V_b^\ell(Y) \right) \quad (4)$$

for all ℓ ; this shows the result. \square

Lemma A.3. *Fix ℓ and ℓ' with $\ell < \ell'$, and let x^ℓ and $x^{\ell'}$, with $b(x^\ell) = b = b(x^{\ell'})$, be the contracts assigned to $s \in S_b$ in the computations of $C^b(A_b^{\ell+1}) = D_b^\ell$ and $C^b(A_b^{\ell'+1}) = D_b^{\ell'}$, respectively. Then, $x^{\ell'} \Gamma^s x^\ell$.*

Proof. The result follows immediately from Lemma A.2, as $A_b^{\ell+1} \subseteq A_b^{\ell'+1} \subseteq X_b$ by construction. \square

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

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

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

Lemma A.4. *The slot-outcome \tilde{Y} is slot-stable.*

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

$$zP^{i(z)}(\varpi(\tilde{Y}))_{i(z)} = Y_{i(z)}, \quad (5)$$

$$z\Pi^s(\varpi(\tilde{Y}))_s = Y_s. \quad (6)$$

Now, by (5) and the fact that Y is the cumulative offer process outcome, we know that z must be proposed in some step $\underline{\ell}$ of the cumulative offer process. We let $\ell \geq \underline{\ell}$ be the first step of the cumulative offer process for which no slot $s' \in S_{b(z)}$ with $s' \triangleright^{b(z)} s$ is assigned z in the computation of $C^{b(z)}(A_{b(z)}^{\ell+1}) = D_{b(z)}^\ell$. (Such a step ℓ must exist because $z \notin Y$.) We let x^ℓ be the contract assigned to s in the computation of $C^{b(z)}(A_{b(z)}^{\ell+1})$. Since $x^\ell \neq z$, we know that $x^\ell \Pi^s z$. But then, we know by Lemma A.3 that for each $\ell' \geq \ell$, the contract $x^{\ell'}$ assigned to s in the computation of $C^{b(z)}(A_{b(z)}^{\ell'+1})$ has (weakly) higher Π^s -priority than x^ℓ , and hence has (strictly) higher Π^s -priority than z : $x^{\ell'} \Gamma^s x^\ell \Pi^s z$. In particular, then, we must have $Y_s \Pi^s z$, contradicting (6). Thus, there cannot be a slot-contract $\langle z; s \rangle$ that slot-blocks \tilde{Y} . \square

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

Lemma A.5. *For each agent $i \in I$, we have $\tilde{Y}_i \tilde{R}^i \tilde{Z}_i$.*

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

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

If $x \neq (\varpi(\tilde{Z}))_{i(x)}$, then we must have $xP^{i(x)}(\varpi(\tilde{Z}))_{i(x)}$ and it follows immediately that $\langle x; s \rangle$ slot-blocks \tilde{Z} . Meanwhile, if $x = (\varpi(\tilde{Z}))_{i(x)}$, then there is some contract of the form $\langle x; s' \rangle \in \tilde{Z}$, with $s' \neq s$. If $s \triangleright^b s'$, 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 $s' \triangleright^b s$, then we let y be the contract assigned to s' in the computation of $C^{b(x)}(A_{b(x)}^{\ell+1})$. As $s' \triangleright^b s$ and x is not assigned to s' in the computation of $C^{b(x)}(A_{b(x)}^{\ell+1})$, we must have $y \Pi^{s'} x$. Moreover, as z is the first contract in $\varpi(\tilde{Z})$ to be rejected, we know that $yR^{i(y)}(\varpi(\tilde{Z}))_{i(y)}$. 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', \dots of slots (in $S_{b(z)}$) 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 A.4, we know that \tilde{Y} is slot-stable; it then follows from Lemma A.5 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 1

The result follows immediately from Lemma 2 and Lemma 3, 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 2

The result is an immediate consequence of Theorem 1: 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 1, 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 $ZR^i\varpi(\tilde{Y}) = Y$ for each $i \in I$. If $Z \neq Y$, then, we have $ZP^i\varpi(\tilde{Y}) = Y$ for some $i \in I$, so that Z is stable and Pareto dominates Y —a contradiction.

Proof of Theorem 3

Stability is immediate from Theorem 1. Meanwhile, to show (group) strategy-proofness, we suppose that there is some set of agents $I' \subset I$ and $\bar{P}^{I'} \neq P^{I'}$ such that

$$\Phi_{\Pi}(\bar{P}^{I'}, P^{-I'})P^i\Phi_{\Pi}(P^{I'}, P^{-I'}) \quad (7)$$

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

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

by (7) and Theorem A.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).⁴⁰

Proof of Theorem 4

We fix an agent i and let $\bar{\Pi}$ be an unambiguous improvement over Π for i . We let \tilde{Y} and \tilde{Y} be the outcomes of the agent-optimal slot-stable mechanism for the agent–slot market under priorities $\bar{\Pi}$ and Π , respectively (holding precedence fixed).

Claim. *We have $\tilde{Y} \tilde{R}^i \tilde{Y}$, that is, agent i weakly prefers \tilde{Y} , the agent-optimal slot-stable slot-outcome under priorities $\bar{\Pi}$, to \tilde{Y} , the agent-optimal slot-stable slot-outcome under priorities Π .*

Proof. We suppose to the contrary that $\tilde{Y} \tilde{P}^i \tilde{Y}$. We then see immediately that \tilde{Y} must not be slot-stable under priorities $\bar{\Pi}$, as otherwise \tilde{Y} could not be the agent-optimal slot-stable slot-outcome.

We let $\langle x; s \rangle = \tilde{Y}_i$, and consider the alternate extended preferences ranking \tilde{P}^i such that $\langle x; s \rangle \tilde{P}^i \emptyset_i$ and $\emptyset_i \tilde{P}^i \langle y; s' \rangle$ for any $\langle y; s' \rangle \neq \langle x; s \rangle$.⁴¹ Now, we show that \tilde{Y} is stable in the agent–slot matching market under extended preferences $(\tilde{P}^i, \tilde{P}^{-i})$ and priorities $\bar{\Pi}$:

- Individual rationality of \tilde{Y} follows from the individual rationality of \tilde{Y} under extended preferences $(\tilde{P}^i, \tilde{P}^{-i})$ and priorities $\bar{\Pi}$.
- Moreover, \tilde{Y} cannot be slot-blocked by a slot contract $\langle y; s' \rangle$ with $i(y) \neq i$, as then $\langle y; s' \rangle$ would slot-block \tilde{Y} under extended preferences $(\tilde{P}^i, \tilde{P}^{-i})$ and priorities $\bar{\Pi}$.
- Finally, \tilde{Y} cannot be slot-blocked by a slot contract $\langle y; s' \rangle$ with $i(y) = i$, because $\emptyset_i \tilde{P}^i \langle y; s' \rangle$ by construction.

We let \tilde{Z} be the agent-optimal slot-stable slot-outcome under extended preferences $(\tilde{P}^i, \tilde{P}^{-i})$ and priorities $\bar{\Pi}$. As $\tilde{Z} \tilde{P}^i \tilde{Y}$, we must have $\tilde{Z}_i = \langle x; s \rangle \tilde{P}^i \tilde{Y}_i$, by the definition of agent-optimality. But then, we see that i can manipulate the agent-optimal slot-stable mechanism under extended preferences \tilde{P}^i and priorities $\bar{\Pi}$ —if i reports his true (extended) preferences, then he receives \tilde{Y}_i , but if he reports \tilde{P}^i instead, he receives \tilde{Z}_i , which he prefers. But this implies that the agent-optimal slot-stable mechanism is not strategy-proof (in the agent–slot market), contradicting Theorem 1 of Hatfield and Kojima (2009).⁴² \square

⁴⁰Theorem 1 of Hatfield and Kojima (2009) implies that the agent–slot matching mechanism that 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.

⁴¹Note that \tilde{P}^i does not formally extend any preference ranking in the original market; we introduce it only to consider a possible manipulation in the agent–slot market.

⁴²See Footnote 40 for details.

The theorem now follows directly: We have $\varpi(\tilde{Y}) = \Phi_{\tilde{\Pi}}(P^I)$ and $\varpi(\tilde{Y}) = \Phi_{\Pi}(P^I)$ by Theorem A.1, so that

$$\Phi_{\tilde{\Pi}}(P^I) = \varpi(\tilde{Y})R^i\varpi(\tilde{Y}) = \Phi_{\Pi}(P^I)$$

by the preceding claim.