

Matching with (Branch-of-Choice) Contracts at United States Military Academy*

Tayfun Sönmez[†] and Tobias B. Switzer[‡]

May 2011

Abstract

Branch selection is a key decision in a cadet's military career. Cadets at USMA can increase their branch priorities at a fraction of slots by extending their service agreement. This real-life matching problem fills an important gap in market design literature. Although priorities fail a key substitutes condition, the agent-optimal stable mechanism is well-defined, and in contrast to the current USMA mechanism it is fair, stable, and strategy-proof. Adoption of this mechanism benefits cadets and the Army. This new application shows that matching with contracts model is practically relevant beyond traditional domains that satisfy the substitutes condition.

JEL Classification Numbers: C78, D63, D78

Keywords: Market Design, Matching with Contracts, Stability, Strategy-Proofness.

1 Introduction

As a response to historically low retention rates among junior officers, a group of economists and senior officers at West Point's Office of Economic and Manpower Analysis (OEMA) designed a program called the *Officer Career Satisfaction Program* (OCSP) in the mid 2000s to boost

*We thank Federico Echenique, John Hatfield and Fuhito Kojima for helpful comments.

[†]Department of Economics, Boston College, Chestnut Hill, MA 02467. E-mail: sonmezt@bc.edu

[‡]United States Air Force, 6th Special Operations Squadron, Hurlburt Field, FL 32544. E-mail: tobias.switzer@gmail.com

career satisfaction and retention.¹ OSCP consists of three incentive programs, each of which carries a “price” of an additional three years of active duty service for a total of eight years for United States Military Academy (USMA) cadets. The program is completely voluntary and cadets can choose any two of the following three incentive programs: (1) graduate school, (2) Branch-of-Choice, and (3) Base-of-Choice. The first incentive program is an option for two years of fully funded graduate school that vests once the officer completes his initial five-year commitment along with his additional three-year commitment. The second incentive program is a contract where the cadet agrees to serve an additional three years in exchange for a specific career specialty, or “branch” in military language. The final incentive program is a contract where the cadet agrees to serve an additional three years in exchange for the base of the cadet’s choice for his first assignment. Across years 2006 to 2009, the Army has gained a total of 17,596 man-years through the OSCP increasing the expected eight-year continuation rate of officers from 47 percent to 67 percent (Wardynski, Lyle, and Colarusso 2010). Slightly more than 60 percent of this gain is due to a carefully designed Branch-of-Choice incentive program that has been implemented through the *USMA mechanism*, an “indirect extension” of the celebrated Gale-Shapley (1962) *agent-optimal stable mechanism*.

The agent-optimal stable mechanism and its extensions have had a profound influence on market design. Versions of this mechanism are the basis of various labor market clearinghouses² and have recently been adopted by school choice programs in Boston and New York City.³ An important factor in this unmatched success of the agent-optimal stable mechanism is that once the above-mentioned market design applications were carefully formulated, the “Gale-Shapley program” was able to deliver much of the rest. By contrast, the introduction of the branch-of-choice incentive makes the *cadet-branch matching problem* analytically more challenging than basic two-sided matching models. Indeed, the novel branch priorities the Army adopted with its incentive program do not satisfy a *substitutes* condition, which was thought to be a necessary condition to apply the Gale-Shapley program when the Army started implementing the OSCP in 2006. That the substitutes condition is not necessary was only recently discovered by Hatfield and Kojima (2008). It is fair to suggest that the USMA mechanism was designed at a time

¹See Wardynski, Lyle and Colarusso (2010) for a comprehensive review.

²See Roth and Peranson (1999) for the best known of these labor market clearinghouses, the medical resident market in the U.S.

³See Abdulkadiroğlu and Sönmez (2003), Abdulkadiroğlu, Pathak and Roth (2005), Abdulkadiroğlu, Pathak, Roth, and Sönmez (2005).

when market designers did not fully understand priority-based indivisible goods allocation in the absence of the substitutes condition. Fortunately, significant progress has been made since then, most notably by Hatfield and Kojima (2008, 2010). Building on these recent advances, we show that the challenges introduced to cadet-branch matching with the branch-of-choice incentive program can be overcome fairly easily without disrupting the core of the current USMA matching system.

Prior to implementation of the OCSP, the assignment of cadets to branches was a relatively straightforward task. Cadets at USMA have a strict priority ranking known as an *order of merit list* (OML) that is based on a weighted average of academic performance, physical fitness test scores, and military performance. The Army had been using the *serial dictatorship* induced by the OML to assign slots at 16 branches to cadets. The highest-priority cadet was assigned his top choice, the second-priority cadet was assigned his top choice among the remaining choices, and so on. This mechanism is the only mechanism that is *Pareto efficient* and *fair* in the sense that a higher-priority cadet would never envy the assignment of a lower-priority cadet (Balinski and Sönmez 1999). Moreover, it is *strategy-proof*, and hence, truthful-preference revelation was always in the cadets' best interests. With the introduction of branch-of-choice incentives, the Department of the Army decided to change branch priorities as outlined in its October 1, 2007 dated memorandum to the USMA superintendent:

Cadets will be branched with the initial objective that the first 75 percent of each branch allocation is assigned by CPR. The DCS, G-1 may adjust the 75 percent target based on the needs of the Army. The remaining branch allocations will be distributed in CPR order to cadets who are willing to extend their ADSO by three years to secure their branch of choice. Once all volunteers are exhausted, remaining branch allocations will be distributed in CPR order.⁴

Along with this change in branch priorities, the nature of the outcome of the cadet-branch matching problem has changed as well, because cadets are assigned branch-price pairs under the new problem. Since the Kelso and Crawford (1982) extension of the agent-optimal stable mechanism allows for wages, it is a natural candidate to consider as a solution to the new

⁴Here CPR stands for cadet performance rank, ADSO stands for active duty service obligation, and the Deputy Chief of Staff, G-1 (DCS, G-1) is responsible for developing policy pertinent to active duty service obligations.

problem. This approach, however, is not fruitful because a substitutes condition that is key in Kelso and Crawford (1982) is not satisfied in the present context. As discussed above, the market design literature has been silent about matching problems when part of the allocation is assigned based on priorities and the remainder based on willingness to pay an additional cost coupled with priorities.

In this paper, we show that while the USMA mechanism introduced with OCSF lacks many of the desirable features of its predecessor, the Army’s design is very close to what we argue is the perfect mechanism for this task. We believe adoption of this strategy-proof *cadet-optimal stable mechanism* will not only increase cadet satisfaction with the branch-of-choice incentive program but also provide the Army with very valuable information on cadet preferences that can be used to estimate the full potential of branch-of-choice incentive in terms of gained man-years. As mentioned in the above memorandum, the fraction of slots that will be assigned strictly based on cadet priorities is somewhat flexible, but given the nature of the game induced by the USMA mechanism the Army has no way of reliably estimating the impact of the change of this parameter under the current mechanism. If the cadet-optimal stable mechanism is adopted, the Army will have a very precise estimate of how many additional man-years can be obtained with a change in parameters of the problem, including branch capacities and the fraction of these capacities reserved for cadets willing to pay the increased cost. But this is not the only reason why the Army benefits by adopting a strategy-proof mechanism. Access to true cadet preferences is important for the Army since officers who reach the top ranks of the military tend to come from Combat Arms branches that are disproportionately occupied by whites, and the relative lack of minorities in these branches has a significant impact on the diversity of the senior leadership (Lim et al. 2009).

The USMA mechanism induces a relatively complicated game, and, as we have emphasized, it does not inherit any of its predecessor’s appealing properties. Of particular concern is that, truth-telling may no longer be an optimal strategy.⁵ Furthermore, an equilibrium outcome of the USMA mechanism can be *inefficient, unstable, unfair, and may penalize cadets for unambiguous improvements in their order-of-merit list standings*. Building on recent advancements in the market design literature pioneered by Hatfield and Milgrom (2005), we propose the *cadet-optimal stable mechanism* as an alternative. In contrast to the USMA mechanism, this direct mechanism is *strategy-proof, stable, fair, and respects improvements* in cadet order-of-merit list

⁵This observation is first made by Switzer (2011).

standings. We fully diagnose the cause of all undesired features of the USMA mechanism by showing that the entire issue emanates from one “shortcut” the mechanism takes. Under the branch-of-choice incentive program, each branch can be obtained either at the default cost or at the increased cost. There are 16 branches⁶ and hence 32 possible branch-cost pairs that can be assigned to each cadet. The cadet-optimal stable mechanism is a direct mechanism where each cadet submits his preference ranking of as many as all 32 pairs.⁷ In contrast, the USMA mechanism directs each cadet to report a ranking of 16 branches while at the same time signing a *branch-of-choice contract* for each branch he is “willing” to pay the increased cost. Using the branch preferences along with the set of signed branch-of-choice contracts, a *proxy-preference relation* over branch-cost pairs is constructed under the USMA mechanism, and a virtually identical procedure with the cadet-optimal stable mechanism is used to determine the final outcome. Hence, while cadets can truthfully declare their preferences under our proposed cadet-optimal stable mechanism, cadets need to represent them as best as they can using the smaller strategy space available under the USMA mechanism. Whereas agents have 32! possible preference rankings under the cadet-optimal stable mechanism,⁸ they only have $16! \times 2^{16}$ strategies under the USMA mechanism. Simply put, the strategy-space provided by the USMA mechanism is not nearly rich enough to reasonably represent cadet preferences.⁹ The difference between the strategy spaces results in the USMA mechanism losing most (if not all) of the desirable properties of its predecessor in an otherwise perfect design.

The cadet-branch matching problem is a special case of the *matching with contracts* model (Hatfield and Milgrom 2005) that subsumes and unifies the Gale and Shapley (1962) college admissions model and the Kelso and Crawford (1982) labor market model, among others. This model has attracted a lot of attention not only because it unifies the above-mentioned models but also because it relates matching markets to auctions and formulates how the Gale-Shapley

⁶Female cadets are ineligible for two branches, but we ignore this constraint as it has no impact on our analysis.

⁷Since participation in the OCSP incentive program is fully voluntary, cadets do not need to report the ranking of increased cost alternatives that are worse than all default cost alternatives. Hence each cadet needs to rank anywhere between 16 and 32 alternatives under the cadet-optimal stable mechanism.

⁸While we will make some natural assumptions shrinking the size of the strategy space, the cadet-optimal stable mechanism is well-defined under all 32! preferences.

⁹The adverse impact of limiting strategy space in simpler school choice environments is well-studied. See Haeringer and Klijn (2010) for a theoretical analysis and Calsamiglia, Haeringer, and Klijn (2010) for experimental evidence.

program relates to fixed-point techniques in lattice theory.¹⁰ Indeed, not only is our model a special case of matching with contracts, we believe it is an important special case.

To make this point, we need to briefly discuss two important recent papers. Echenique (2011) has recently shown that under the substitutes condition of Hatfield and Milgrom, the matching with contracts model can be embedded within the Kelso and Crawford (1982) labor market model.¹¹ Many of the results in Hatfield and Milgrom assume the substitutes condition, and thus Echenique (2011) has shown a (highly non-trivial) isomorphism between these results and their counterparts in Kelso and Crawford (1982).¹² Echenique (2011) emphasizes that substitutability is key for the embedding to work. In particular, he indicates that a recent theory paper by Hatfield and Kojima (2010) analyzes matching with contracts under weaker conditions, and his embedding does not work under their conditions. Although Hatfield and Kojima (2010) do not offer any applications under their weaker *unilateral substitutes* condition, they show that many of the key results still persist under this condition. Remarkably, although the substitutability condition fails in the context of cadet-branch matching, the unilateral substitutes condition is satisfied. And thanks to Hatfield and Kojima (2010), results of Hatfield and Milgrom (2005) that are relevant for cadet-branch matching can all be extended under the unilateral substitutes condition. In this sense, our paper presents the first practical application of matching with contracts that cannot be embedded into the Kelso and Crawford labor market model.

While our main focus is solving the Army’s cadet-branch matching problem, we also extend our basic model to analyze priority-based allocation of indivisible goods, when agents can “buy” priority for some of the slots. We allow for a variety of costs for each object type and show how the Hatfield-Kojima-Milgrom approach can be used to solve these problems. We present an application to centralized school admissions where part of the seats are allocated at low-tuition based on priorities while the rest are reserved for high-tuition students. This is reminiscent of high-tuition seats being reserved for out-of-state students at state universities, although our model is relevant for centralized admissions. Another application is parking space assignment where allocation is based on priorities although priorities for some slots can be improved through higher-cost permits.

¹⁰Recent contributions in this research program include Hatfield and Kojima (2008, 2010), Ostrovsky (2008), and Westkamp (2010).

¹¹Kelso and Crawford (1982) build on the analysis of Crawford and Knoer (1981).

¹²See also Kominers (2011) for an extension to many-to-many matching.

The rest of the paper is organized as follows. In Section 2, we introduce the cadet-branch matching problem as well as the USMA mechanism. In Section 3, we introduce the cadet-optimal stable mechanism and show that it satisfies a number of desirable properties. In Section 4, we show that the USMA mechanism is a rough approximation of the cadet-optimal stable mechanism and as a result it fails all of the above-mentioned desirable properties. In Section 5, we extend the model and present applications to parking space assignment and school admissions. We conclude in Section 6.

2 The Cadet-Branch Matching Problem

Prior to 2006, the assignment of cadets to branches at USMA was a basic application of a *student placement problem* (Balinski and Sönmez 1999): A number of cadets were to be assigned to a number of branches based on their preferences over branches, capacities of the branches, and a unique priority ranking of the cadets. This priority ranking of cadets is known as *order-of-merit list* (OML) in military language, and it is based on a *order-of-merit score* (OMS) that is constructed as a weighted average of academic performance, physical fitness test scores, and military performance. The outcome of a problem is a *matching* of cadets to branches such that no cadet is assigned to more than one branch and no branch is assigned more cadets than its capacity. When assigning cadets to branches, it is important to follow the OML: A matching is *fair* if a cadet never envies a lower-priority cadet. When there is a unique priority ranking of agents, there is a perfect solution to a student placement problem. Simply assign the slots at branches to cadets one at a time based on their preferences following their unique ranking: The cadet with the highest priority receives his top choice, the next cadet receives his top choice among the remaining choices, and so on. This direct mechanism is known as a *simple serial dictatorship*, and it is the only mechanism that is *fair* and *Pareto efficient* (Balinski and Sönmez 1999). It also has several other plausible properties, including *strategy-proofness*, and it was the mechanism of choice at USMA prior to 2006.

2.1 The Model

Each cadet is committed to serve in active military service for five years upon completing his studies, although the Army naturally wishes them to serve for longer periods. To encourage

cadets to serve longer, the Army changed the nature of the cadet-branch matching in 2006. Cadets are now given an option to sign one or more *branch-of-choice* contracts which increase their priorities at branches of their choosing in exchange for three additional years of active military service. A branch-of-choice contract does not guarantee a cadet will receive a slot at a branch and nor does signing one necessarily oblige him for the three additional years of service even if the cadet is assigned a slot at the branch. Loosely speaking, the Army has changed the priorities only for the last 25 percent of the slots at each branch and has given priority to cadets who have committed for three additional years of service at these slots. A cadet who signed a branch-of-choice contract is obliged to serve the additional three years of service only if he receives a slot from the last 25 percent of the capacity of a branch.

The addition of branch-of-choice contracts changes the structure of the cadet-branch matching problem considerably. Most importantly, the outcome of the problem is no longer merely an assignment of branches to cadets but rather an assignment of branches along with the terms of these assignments. Cadets now have to evaluate branches together with their “cost” (in terms of service years), and the above-mentioned fairness requirement needs to be modified to take the branch-of-choice contracts into consideration. We are now ready to formally introduce the problem.

A **cadet-branch matching problem** consists of

1. a finite set of cadets $I = \{i_1, i_2, \dots, i_n\}$,
2. a finite set of branches $B = \{b_1, b_2, \dots, b_m\}$,
3. a vector of branch capacities $q = (q_b)_{b \in B}$ with $\sum_{b \in B} q_b \geq n$,
4. a pair of “terms” $T = \{t_0, t_+\}$,
5. a list of cadet preferences $P = (P_{i_1}, P_{i_2}, \dots, P_{i_n})$ over $B \times T$, and
6. a priority ranking $\pi : I \rightarrow \{1, \dots, n\}$.

Since the problem is motivated by the application at military academy, we refer t_0 as the *base cost* and t_+ as the *increased cost*. We assume that cadet preferences are strict and they are such that for any cadet $i \in I$ and any pair of branches b, b' ,

$$(b, t_0)P_i(b', t_0) \Leftrightarrow (b, t_+)P_i(b', t_+).$$

Hence cadet preferences over branches is independent of the service obligation, and thus each cadet has well-defined preferences over branches. Let \succ_i denote cadet preferences over branches alone. For any cadet i and any pair of branches b, b' , we have

$$b \succ_i b' \Leftrightarrow (b, t_0)P_i(b', t_0) \Leftrightarrow (b, t_+)P_i(b', t_+).$$

Let \mathcal{P} denote the set of all preferences over $B \times T$, and \mathcal{Q} denote the set of all preferences over B .

A **contract** $x = (i, b, t) \in I \times B \times T$ specifies a cadet i , a branch b and the terms of their match. Let $X \equiv I \times B \times T$ be the set of all contracts. A **branch-of-choice contract** $x = (i, b, t)$ is a contract with increased cost $t = t_+$. Given a contract $x = (i, b, t)$, let $x_I = i$ denote the cadet, $x_B = b$ denote the branch, and $x_T = t$ denote the terms of the contract x .

An **allocation** $X' \subset X$ is a set of contracts such that each cadet appears in only one contract and no branch appears in more contracts than its capacity. Let \mathcal{X} denote the set of all allocations.

Given a cadet $i \in I$ and an allocation $X' \subset X$ with $(i, b, t) \in X'$, let $X'(i) = (b, t)$ denote the **assignment** of cadet i under allocation X' .

For a given problem, an allocation X' is **fair** if $X'(j)P_iX'(i) \Rightarrow \pi(j) < \pi(i)$ for any pair of cadets i, j . That is, a higher-priority cadet can never envy the assignment of a lower-priority cadet under a fair allocation. Note that it is still possible for a higher-priority cadet to envy the branch assigned to a lower-priority cadet under a fair allocation. Consider a high-priority cadet i with an assignment $X'(i) = (b(i), t_0)$ and a low-priority cadet j with $X'(j) = (b(j), t_+)$. While fairness rules out $X'(j)P_iX'(i)$, it is still possible that $b(j) \succ_i b(i)$. A low-priority cadet may be able to get a more preferred branch, because he is willing to pay a higher price for it.

2.2 The USMA Mechanism

A **mechanism** is a **strategy space** S_i for each cadet i along with an **outcome function** $\varphi : (S_1 \times S_2 \times \cdots \times S_n) \rightarrow \mathcal{X}$ that selects an allocation for each strategy vector $(s_1, s_2, \dots, s_n) \in (S_1 \times S_2 \times \cdots \times S_n)$. Given a cadet i and strategy a profile $s \in S$, let s_{-i} denote the strategy of all cadets except cadet i .

A **direct mechanism** is a mechanism where the strategy space is simply the set of preferences \mathcal{P} for each cadet i . Hence a direct mechanism is simply a function $\psi : \mathcal{P}^n \rightarrow \mathcal{X}$ that selects an allocation for each preference profile.

The strategy space of each cadet is $\mathcal{Q} \times 2^B$ under the current USMA mechanism. That is, each cadet is asked to choose

1. a ranking of branches alone, and
2. a number of branches (possibly none) for which the cadet is asked to sign a branch-of-choice contract.

Let (\succ'_i, B_i) be the strategy choice of cadet i under the USMA mechanism for a given problem. Here \succ'_i is the reported preferences of cadet i over branches and B_i is the set of branches for which cadet i has signed a branch-of-choice contract. The branch-of-choices chosen by cadet i indicates that cadet i is willing to pay the additional cost of t_+ for each branch for which he signs a branch-of-choice contract. The cadet will need to pay the additional cost only if he receives one of the last 25 percent of the slots for which he is favored.

Under the USMA mechanism, cadets who sign a branch-of-choice contract for branch b receive higher priority for the last λ fraction of slots at branch b (where $\lambda = 0.25$ at the USMA). Hence, while the priority for the first $(1 - \lambda)q_b$ slots is strictly based on the merit ranking π , the **adjusted priority ranking** π_b^+ for the last λq_b slots is constructed in the following natural way:¹³

- $\pi_b^+(i) < \pi_b^+(j)$ for any two cadets i, j such that $b \in B_i$ and $b \notin B_j$,
- $\pi_b^+(i) < \pi_b^+(j)$ for any two cadets i, j such that $\pi(i) < \pi(j)$, $b \in B_i$ and $b \in B_j$,
- $\pi_b^+(i) < \pi_b^+(j)$ for any two cadets i, j such that $\pi(i) < \pi(j)$, $b \notin B_i$ and $b \notin B_j$.

So the relative priority of two cadets does not change under π_b^+ unless one of the cadets has signed a branch-of-choice contract for branch b while the other has not. For any two such cadets, the cadet who has signed a branch-of-choice contract has higher priority under π_b^+ than the one who has not signed one for the last λq_b slots at branch b .

We are ready to describe the algorithm USMA uses to assign branch-price pairs to cadets given their strategy choices $(\succ'_i, B_i)_{i \in I}$. For any branch b , we construct the adjusted priorities π_b^+ as described above.

¹³Throughout the paper we will assume that λq_b and $(1 - \lambda)q_b$ as are integer values for any branch b . This is for expositional simplicity and has no impact on the analysis. One can easily use rounded values for the entire analysis.

The USMA Algorithm

Step 1: Each cadet i applies to his top-choice branch under his submitted preferences \succ'_i . Each branch b tentatively accepts the top $(1 - \lambda)q_b$ candidates based on the merit ranking π . Among the remaining applicants it tentatively accepts the top λq_b candidates based on the adjusted priorities π^+ . Any remaining applicants are rejected.

Step 2: Each cadet i who is rejected at Step 1 applies to his next-choice branch under his submitted preferences \succ'_i . For the first $(1 - \lambda)q_b$ slots, each branch b reviews the new applicants along with those held from Step 1, and tentatively accepts the top $(1 - \lambda)q_b$ based on the merit ranking π . For the last λq_b slots, branch b considers all remaining applicants except those already rejected by branch b at Step 1 and tentatively accepts the top λq_b of them based on the adjusted priorities π^+ .¹⁴ Any remaining applicants are rejected.

In general, at

Step k : Each cadet i who is rejected at Step $(k-1)$ applies to his next-choice branch under his submitted preferences \succ'_i . For the first $(1 - \lambda)q_b$ slots, each branch b reviews the new applicants along with those held from Step $(k-1)$, and tentatively accepts the top $(1 - \lambda)q_b$ based on the merit ranking π . For the last λq_b slots, branch b considers all remaining applicants except those already rejected in earlier steps and tentatively accepts the top λq_b of them based on the adjusted priorities π^+ . Any remaining applicants are rejected.

The algorithm terminates when no applicant is rejected. All tentative assignments are finalized at that point. For any branch b ,

1. any cadet who is assigned one of the top $(1 - \lambda)q_b$ slots by the algorithm is charged the lower cost t_0 ,
2. any cadet who is assigned one of the last λq_b slots is charged the higher cost t_+ if he has signed a branch-of-choice contract for branch b , and
3. any cadet who is assigned one of the last λq_b slots is charged the lower cost t_0 if he has not signed a branch-of-choice contract for branch b .

Given a strategy profile $s = (\succ'_i, B_i)_{i \in I}$ and an order-of-merit list π , let $\varphi_\pi^{WP}(s)$ denote the outcome of the USMA mechanism. Whenever the order-of-merit list π is fixed throughout an analysis, we will suppress π and denote the outcome as $\varphi^{WP}(s)$.

¹⁴Observe that a cadet can be tentatively accepted for a branch b for the first $(1 - \lambda)q_b$ slots in Step 1 but not in Step 2, although he is tentatively accepted for the last λq_b slots in Step 2.

The USMA algorithm reduces to the simple serial dictatorship induced by merit ranking π when $\lambda = 0$. For the case of $\lambda = 0$, the USMA algorithm can also be interpreted as a special case of the celebrated **agent-proposing deferred acceptance algorithm** (Gale and Shapley 1962), which allows for a different priority ranking at each branch. Both of these mechanisms are very well behaved: Not only do they always result in a fair allocation, but truthful preference revelation is a dominant strategy for all cadets under either mechanism.

The analysis of the USMA mechanism for $\lambda > 0$ is somewhat more delicate. That is because not only may truthful preference revelation be suboptimal under the USMA mechanism, but also the optimal choice of branch-of-choice contracts is a challenging task. Before we analyze the equilibrium outcomes of the USMA mechanism, we will show that the cadet-branch matching problem is a special case of a recent influential model by Hatfield and Milgrom (2005). That will not only allow us to propose an alternative mechanism to the cadet-branch matching problem but also will allow us to better understand the equilibrium outcomes of the USMA mechanism.

3 Matching with Contracts

The cadet-branch matching problem is a special case of the **matching with contracts** model (Hatfield and Milgrom 2005) that subsumes and unifies the Gale and Shapley (1962) college admissions model and the Kelso and Crawford (1982) labor market model, among others. In the original Hatfield-Milgrom model, each branch (hospitals in their framework) has preferences over sets of agent-cost pairs. These hospital preferences induce a *chosen set* from each set of contracts, and it is this chosen set (rather than hospital preferences) that is key in the model. In the present framework, branches are not agents and they do not have preferences. However, branches have priorities over cadet-cost pairs, and these priorities also induce chosen sets. This is the sense in which the cadet-branch matching problem is a special case of matching with contracts. We next present a few key concepts from matching with contracts before we propose a new mechanism for cadet-branch matching. Recall that each cadet i has preferences P_i over all branch-cost pairs. Equivalently, each cadet i has strict preferences over all contracts that include him. Since no cadet can be “forced” to pay the increased cost, the relevant contracts with increased cost t_+ are those that are more preferred to at least one contract with the base cost t_0 .

Definition: A contract (i, b, t_+) is **unacceptable** for cadet i if $(b', t_0) P_i (b, t_+)$ for all $b' \in B$.

Given a set of contracts X' , let the **chosen set of cadet i** from X' , $C_i(X')$, be \emptyset if all contracts in X' that include cadet i are unacceptable and the singleton set consisting of the most preferred contract including cadet i in X' under P_i otherwise. That is,

$$C_i(X') = \left\{ (i, b, t) \in X' : (b, t) P_i(b', t') \text{ for any } (i, b', t') \in X' \setminus \{(i, b, t)\} \right\}$$

unless all contracts in X' that include cadet i are unacceptable.

In general, the chosen set of branch b from a set of contracts X' depends on the policy on who has higher claims for slots in branch b . Here our focus is the current USMA policy, where cadets with higher OML ranking have higher claims for the top $(1 - \lambda)q_b$ slots and the priority for the last λq_b slots is adjusted to favor cadets who are willing to pay the increased service cost. We are ready to define the **USMA chosen set of branch b** from a set of contracts X' . For a given order-of-merit priority ranking π and $\lambda \in [0, 1]$, order all contracts in X' based on the merit score of the cadet and construct USMA chosen set $C_b(X')$ and the **USMA rejected set** $R_b(X') = X' \setminus C_b(X')$ as follows:

Phase 0: Remove all contracts that involve another branch b' and add them all to rejected set $R_b(X')$. Hence each contract that survives Phase 0 involves branch b .

Phase 1: For the first $(1 - \lambda)q_b$ potential elements of $C_b(X')$, simply choose the contracts with highest-priority cadets one at a time. When two contracts of the same cadet are available, choose the contract with the base cost t_0 and reject the other one, including it in $R_b(X')$. Continue until either all contracts are considered or $(1 - \lambda)q_b$ elements are chosen for $C_b(X')$. If the former happens, terminate the procedure and if the latter happens proceed with Phase 2.1.

Phase 2.1: For the last λq_b potential elements of $C_b(X')$, we give priority to contracts with increased cost t_+ . Hence in this phase only consider branch-of-choice contracts (i.e. contracts with increased cost t_+) and among them include in $C_b(X')$ the contracts with highest-priority cadets. If any cadet covered in Phase 2.1 has two contracts in X' reject the contract with the base cost t_0 , including it in $R_b(X')$. Continue until either all branch-of-choice contracts are considered in X' or $C_b(X')$ fills all q_b elements. For the latter case, reject all remaining contracts, including them in $R_b(X')$ and terminate the procedure. For the former case, terminate the procedure if all contracts in X' are considered and proceed with the Phase 2.2 otherwise.

Phase 2.2: By construction, all remaining contracts in X' have the base cost t_0 . Include in $C_b(X')$ the contracts with highest-priority cadets one at a time until either all contracts in X'

are considered or $C_b(X')$ fills all q_b elements. Reject any remaining contracts, including them in $R_b(X')$.

3.1 Stability

Since the seminal paper of Gale and Shapley (1962), a condition known as stability has been central to the analysis of two-sided matching markets as well as allocation of indivisible goods based on priorities.¹⁵ Formally, an allocation X' is **stable** if

1. $\bigcup_{i \in I} C_i(X') = X'$,
2. $\bigcup_{b \in B} C_b(X') = X'$, and
3. there exists no cadet i , branch b , and contract $x = (i, b, t) \in X \setminus X'$ such that

$$\underbrace{\{x\} = C_i(X' \cup \{x\})}_{\left(\text{i.e. } (b,t)P_i X'(i)\right)} \text{ and } x \in C_b(X' \cup \{x\}).$$

In the context of cadet-branch matching, the only plausible allocations are the stable ones. Note that if the first requirement fails then there is a cadet who prefers to reject a contract that involves him (or equivalently, there is a cadet who is given an unacceptable contract); if the second requirement fails then there exists a branch that would rather reject one of the contracts that includes it; and if the third requirement fails then there exists an unselected contract (i, b, t) where not only cadet i prefers pair (b, t) to his assignment, but also contract x has sufficiently high priority to be selected by branch b given its composition.

3.2 Substitutes, Unilateral Substitutes, and the Law of Aggregate Demand

The following two properties of branch priorities have played an important role in the analysis of matching with contracts:

Definition: Priorities satisfy the **law of aggregate demand** for branch b if $X' \subset X'' \Rightarrow |C_b(X')| \leq |C_b(X'')|$.

¹⁵See Roth and Sotomayor (1990) and Sönmez and Ünver (2010) for comprehensive surveys on the role of stability in two-sided matching markets and allocation of indivisible goods based on priorities.

That is, the size of the chosen set never shrinks as the set of contracts grows under the law of aggregate demand.

Lemma 1: The USMA priorities satisfy the law of aggregate demand for each branch b .

Definition: Elements of X are **substitutes** for branch b if for all $X' \subset X'' \subseteq X$ we have $R_b(X') \subseteq R_b(X'')$.

That is, contracts are substitutes if a contract that is chosen from a larger set X'' is chosen from any of its subsets $X' \subset X''$ as well. Equivalently, any contract that is rejected from a smaller set X' is also rejected from any larger set X'' that contains X' . If elements of X are substitutes, then the set of stable allocations is non-empty (Hatfield and Milgrom 2005). Indeed, until Hatfield and Kojima (2008) showed otherwise, it was thought to be a necessary condition for the guaranteed existence of a stable allocation. The following example shows that contracts are not necessarily substitutes under the USMA priorities.

Example 1: Let $I = \{i_1, i_2, i_3\}$, $B = \{b_1, b_2\}$ with $q_{b_1} = q_{b_2} = 2$. Let the merit ranking π order cadets as i_1, i_2, i_3 and $\lambda = 0.5$. Let $X' = \{(i_2, b_1, t_0), (i_2, b_1, t_+)\}$ and $X'' = \{(i_1, b_1, t_0), (i_2, b_1, t_0), (i_2, b_1, t_+)\}$. We have $C_{b_1}(X') = \{(i_2, b_1, t_0)\}$, $R_{b_1}(X') = \{(i_2, b_1, t_+)\}$, $C_{b_1}(X'') = \{(i_1, b_1, t_0), (i_2, b_1, t_+)\}$ and $R_{b_1}(X'') = \{(i_2, b_1, t_0)\}$. Hence, even though contract (i_2, b_1, t_+) is not rejected from X'' , it is rejected from X' , which is included in X'' . \diamond

Consider two contracts $(i, b, t_0), (i, b, t_+)$ *including the same cadet*: Under the USMA chosen set C_b , contract (i, b, t_0) might be rejected at the expense of (i, b, t_+) from a larger set X'' , while the choice is reversed for a subset X' of X'' . That is the only reason why a contract can be rejected from a smaller set despite being chosen from a larger one. Therefore, the USMA chosen set C_b satisfies the following weaker condition, introduced by Hatfield and Kojima (2010):

Definition: Elements of X are **unilateral substitutes** for branch b if, whenever a contract $x = (i, b, t)$ is rejected from a smaller set X' *even though x is the only contract in X' that includes cadet i* , contract x is also rejected from a larger set X'' that includes X' .

Note that the emphasized part of the definition is what differentiates unilateral substitutes condition from substitutes, and since any “reversal” in the choice of two contracts involves the same cadet, the USMA priorities satisfy the unilateral substitutes condition. It turns out that, this will be key in the context of cadet-branch matching.

Lemma 2: Elements of X are unilateral substitutes for each branch b under the USMA priorities.

3.3 Cumulative Offer Algorithm and the Cadet-Optimal Stable Mechanism

We are ready to introduce the **cadet-optimal stable mechanism (cosm)**, which is simply the extension of the celebrated **agent-optimal stable mechanism** (Gale and Shapley 1962), to matching with contracts:

The strategy space of each cadet is \mathcal{P} under the cosm, and hence it is a direct mechanism. Given a preference profile $P \in \mathcal{P}$, the following algorithm can be used to find the outcome of cosm.

Cumulative Offer Algorithm¹⁶

Step 1: Start the offer process with the highest-merit-score cadet $\pi(1) = i(1)$. Cadet $i(1)$ offers his first-choice contract $x_1 = (i(1), b(1), t)$ to branch $b(1)$ that is involved in this contract. Branch $b(1)$ holds the contract if $x_1 \in C_{b(1)}(\{x_1\})$ and rejects it otherwise. Let $A_{b(1)}(1) = \{x_1\}$ and $A_b(1) = \emptyset$ for all $b \in B \setminus \{b(1)\}$.

Step 2: Let $i(2)$ be the highest-merit-score cadet for whom no contract is currently held by any branch. Cadet $i(2)$ offers his most-preferred contract $x_2 = (i(2), b(2), t)$ that has not been rejected in the previous step to branch $b(2)$. Branch $b(2)$ holds the contract if $x_2 \in C_{b(2)}(A_{b(2)}(1) \cup \{x_2\})$ and rejects it otherwise. Let $A_{b(2)}(2) = A_{b(2)}(1) \cup \{x_2\}$ and $A_b(2) = A_b(1)$ for all $b \in B \setminus \{b(2)\}$.

In general, at

Step k: Let $i(k)$ be the highest-merit-score cadet for whom no contract is currently held by any branch. Cadet $i(k)$ offers his most-preferred contract $x_k = (i(k), b(k), t)$ that has not been rejected in previous steps to branch $b(k)$. Branch $b(k)$ holds the contract if $x_k \in C_{b(k)}(A_{b(k)}(k-1) \cup \{x_k\})$ and rejects it otherwise. Let $A_{b(k)}(k) = A_{b(k)}(k-1) \cup \{x_k\}$ and $A_b(k) = A_b(k-1)$ for all $b \in B \setminus \{b(k-1)\}$.

The algorithm terminates when all cadets have an offer that is on hold by a branch. Since there are a finite number of contracts, the algorithm terminates after a finite number T of steps. All contracts held at this final Step T are finalized and the final allocation is $\bigcup_{b \in B} C_b(A_T)$.

¹⁶The following description is borrowed mostly from Hatfield and Kojima (2010).

Remark 1: While the choice of the cadet making the offer at any given step is uniquely defined in the above-described cumulative offer algorithm, the same outcome is obtained regardless of the choice of cadet as long as there is no contract held by any branch that involves the cadet making the offer. Indeed, Hatfield and Kojima (2010) describe the algorithm without explicitly specifying the order of agents making offers.

Given a preference profile $P \in \mathcal{P}$ and order-of-merit list π , let $\varphi_\pi^{GS}(P)$ denote the outcome of the (Gale-Shapley) cadet-optimal stable mechanism. Whenever π is fixed throughout an analysis, we will suppress it and denote the outcome as $\varphi^{GS}(P)$.

The first result we present justifies the name of the cadet-optimal stable mechanism.

Proposition 1: The cumulative offer algorithm produces a stable allocation under the USMA branch priorities. Moreover, this allocation is weakly preferred by any cadet to any stable allocation.

Given Proposition 1, it is tempting to conclude that the fairness of the cadet-optimal stable mechanism immediately follows as a corollary. However, stability (even under USMA priorities) does not imply fairness in our current setup. Note that we have not assumed that a cadet necessarily prefers a pair (b, t_0) to pair (b, t_+) . While this assumption is very natural, our entire analysis is valid even in its absence. And there may be cadets who might prefer a longer-term commitment for various reasons. That is one reason why a stable allocation might not be fair.

Example 2. There are two cadets i_1, i_2 , one branch b_1 with two slots, and $\lambda = 0.5$. Cadet i_1 has higher priority than cadet i_2 , and their preferences are as follows:

$$(b_1, t_+)P_{i_1}(b_1, t_0) \text{ and } (b_1, t_0)P_{i_2}(b_1, t_+).$$

Consider the allocation $X' = \{(i_1, b_1, t_0), (i_2, b_2, t_+)\}$. Because of his unusual preferences, the higher-priority cadet i_1 envies the assignment of the lower-priority cadet i_2 . Therefore, X' is not fair. Nevertheless, X' is stable because branch b gives priority to contract (i_1, b_1, t_0) over contract (i_1, b_1, t_+) for the first slot. \diamond

Definition: A mechanism is stable if it always chooses a **stable** allocation. Similarly, a mechanism is **fair** if it always chooses a fair allocation.

We next show that the cadet-optimal stable mechanism is fair, even though not all stable mechanisms need to be fair, as Example 2 shows.

Proposition 2: The cadet-optimal stable mechanism is fair under USMA branch priorities.

One of the remarkable properties of the cadet-optimal stable mechanism is that it is always in cadets' best interests to be entirely truthful about their preferences. Hence, cadets can never benefit from “gaming” this mechanism.

Definition: A direct mechanism φ is **strategy-proof** if

$$\varphi(P_i, P_{-i})P_i\varphi(P'_i, P_{-i}) \text{ for any } i \in I, P_{-i} \in \mathcal{P}^{n-1} \text{ and } P_i, P'_i \in \mathcal{P}.$$

That is, no matter which cadet we consider, no matter what his true preferences P_i are, no matter what other preferences P_{-i} other cadets report (true or not), and no matter which potential “misrepresentation” P'_i cadet i considers, truthful preference revelation is in his best interests.

Proposition 3: The cadet-optimal stable mechanism is strategy-proof under USMA branch priorities.

Remark 2: We rely on the unilateral substitutes condition to prove Propositions 1 and 3. An alternative path is to show that USMA priorities are *substitutable completable*, a condition recently introduced by Hatfield and Kominers (2011), and to use Lemma 16, and Theorems 17 and 18 in their paper.

One of the most important parameters of the cadet-branch matching problem is the order-of-merit list π . And clearly, a reasonable mechanism would never penalize a cadet as a result of an improvement of his standing in the order-of-merit list. Given two order-of-merit lists π_1, π_2 , we will say that π_1 is an **unambiguous improvement** for cadet i over π_2 if

1. the relative ranking between all cadets except cadet i remains exactly the same between π_1 and π_2 , although
2. the standing of cadet i is strictly better under π_1 than under π_2 .

Definition: A mechanism **respects improvements** if a cadet never receives a strictly worse assignment as a result of an unambiguous improvement of his priority ranking.

Proposition 4: The cadet-optimal stable mechanism respects improvements under USMA branch priorities.

4 The Case Against the USMA Mechanism

Analyzing the cadet-optimal stable mechanism is rather easy; not only it is a direct mechanism, it is also strategy-proof, so there is good reason to believe that submitted cadet preferences are truthful. Indeed, vulnerability of school choice mechanisms to preference manipulation has been the main driving force behind the recent school choice reforms in Boston, Chicago, and throughout England. (See Abdulkadiroğlu and Sönmez 2003, Abdulkadiroğlu, Pathak, Roth and Sönmez 2005, Kesten 2010, Pathak and Sönmez 2008, 2011.)

The USMA mechanism, on the other hand, is not a direct mechanism. Under the USMA mechanism, each cadet submits his preferences over individual branches, and by signing branch-of-choice contracts he indicates for which of these branches he is willing to pay the additional cost of t_+ for an increased priority at the last 25 percent of the slots. Since the USMA mechanism is based on a strategy-proof mechanism, it is natural to ask whether truthful preference revelation over branches is still in the best interests of the cadets. As we show in the next example, this may no longer be the case.

Example 3 (Truthful preference revelation may be suboptimal under USMA mechanism): There are 3 cadets i_1, i_2, i_3 and 3 branches b_1, b_2, b_3 with one slot each. Cadets are ordered as i_1, i_2, i_3 under the order-of-merit list π and $\lambda = 1$, so any cadet who has signed a branch-of-choice contract is given priority at each branch.¹⁷

Suppose that the true preferences of cadet i_3 over branch-price pairs are as follows:

$$(b_1, t_0)P_{i_3}(b_2, t_0)P_{i_3}(b_1, t_+)P_{i_3}(b_2, t_+)P_{i_3}(b_3, t_0)$$

First, consider the situation where the submitted strategies of cadets i_1 and i_2 are as follows:

$$\begin{aligned} \tilde{s}_{i_1} &= (\tilde{\succ}_{i_1}, \tilde{B}_{i_1}) && \text{with } b_1 \tilde{\succ}_{i_1} b_3 \tilde{\succ}_{i_1} b_2 \text{ and } \tilde{B}_{i_1} = \emptyset, \\ \tilde{s}_{i_2} &= (\tilde{\succ}_{i_2}, \tilde{B}_{i_2}) && \text{with } b_1 \tilde{\succ}_{i_2} b_3 \tilde{\succ}_{i_2} b_2 \text{ and } \tilde{B}_{i_2} = \emptyset. \end{aligned}$$

Given $(\tilde{s}_{i_1}, \tilde{s}_{i_2})$, the best possible outcome for the lowest-priority cadet i_3 is (b_2, t_0) , which he

¹⁷It is straightforward to extend this example for $\lambda = 0.25$, but that would require an example with 9 additional agents and 3 additional slots at each branch.

can obtain with any one of the following three strategies:

$$\begin{aligned}
s_{i_3}^1 &= (\succ_{i_3}, B_{i_3}^1) && \text{with } b_1 \succ_{i_3} b_2 \succ_{i_3} b_3 \text{ and } B_{i_3}^1 = \emptyset, \\
s_{i_3}^2 &= (\succ'_{i_3}, B_{i_3}^1) && \text{with } b_2 \succ'_{i_3} b_1 \succ'_{i_3} b_3 \text{ and } B_{i_3}^1 = \emptyset, \\
s_{i_3}^3 &= (\succ'_{i_3}, B_{i_3}^2) && \text{with } b_2 \succ'_{i_3} b_1 \succ'_{i_3} b_3 \text{ and } B_{i_3}^2 = \{b_1\}.
\end{aligned}$$

Hence, the best response of cadet i_3 to $(\tilde{s}_{i_1}, \tilde{s}_{i_2})$ is $\{s_{i_3}^1, s_{i_3}^2, s_{i_3}^3\}$. Therefore, the only two strategies that can possibly dominate $s_{i_3}^3$ are $s_{i_3}^1$ or $s_{i_3}^2$. We next show that neither strategy dominates $s_{i_3}^3$. Consider the situation where the submitted strategies of cadets i_1 and i_2 are as follows:

$$\begin{aligned}
\hat{s}_{i_1} &= (\hat{\succ}_{i_1}, \hat{B}_{i_1}) && \text{with } b_2 \hat{\succ}_{i_1} b_1 \hat{\succ}_{i_1} b_3 \text{ and } \hat{B}_{i_1} = \emptyset, \\
\hat{s}_{i_2} &= (\hat{\succ}_{i_2}, \hat{B}_{i_2}) && \text{with } b_1 \hat{\succ}_{i_2} b_2 \hat{\succ}_{i_2} b_3 \text{ and } \hat{B}_{i_2} = \{b_2\}.
\end{aligned}$$

Given $(\hat{s}_{i_1}, \hat{s}_{i_2})$, the best possible outcome for cadet i_3 is (b_1, t_+) . While cadet i_3 can obtain this outcome upon submitting $s_{i_3}^3$, he will receive the inferior outcome (b_3, t_0) upon submitting either $s_{i_3}^1$ or $s_{i_3}^2$. Hence, neither strategy dominates $s_{i_3}^3$, showing that it is not a dominated strategy. Hence misrepresenting preferences over branches might be optimal under the USMA mechanism. \diamond

Example 3 shows that one cannot assume that submitted cadet preferences over branches are necessarily truthful under the USMA mechanism. This makes analysis of the USMA mechanism more challenging than the cosm. For the rest of this section, we will consider the Nash equilibrium outcomes of the USMA mechanism.

Definition: Given a mechanism (S, φ) , the strategy profile $s \in S$ is a **Nash equilibrium** if $\varphi(s)P_i\varphi(s_{-i}, s'_i)$ for any cadet i and any alternative strategy $s'_i \in S_i$ cadet i can choose.

We next show that the USMA mechanism can be interpreted as an ‘‘approximation’’ of the cadet-optimal stable mechanism. Fix a cadet-branch problem and let $s_i = (\succ'_i, B_i)$ be the strategy choice of cadet i under the USMA mechanism. Construct the **proxy preference relation** $P_i^* = P_i(\succ'_i, B_i)$ as follows:

1. $b \succ'_i b' \Rightarrow (b, t_0) P_i^* (b', t_0)$ for any pair of branches b, b' ,
2. $(b, t_0) P_i^* (b, t_+)$ for any branch $b \in B_i$, and moreover there exists no branch $b' \in B$ with $(b, t_0) P_i^* (b', t_0) P_i^* (b, t_+)$, and

3. $(b, t_0) P_i^* (b', t_+)$ for any $b \in B$ and $b' \notin B_i$.

In other words, the proxy preference relation P_i^* extends cadet i 's submitted preferences \succ'_i over branches to branch-price pairs by simply inserting the (b, t_+) option right after (b, t_0) for each branch b for which cadet i has signed a branch-of-choice contract.

Given a cadet-branch matching problem, let s be a Nash equilibrium strategy profile under the USMA mechanism and P^* be the proxy preferences as described above. The equilibrium outcome $\varphi^{WP}(s)$ is equal to the outcome of the cadet-optimal stable mechanism under the proxy preference profile P^* .

Proposition 5: Let $s = (\succ_i, B_i)_{i \in I}$ be a Nash equilibrium strategy profile under the USMA mechanism and $P^* = (P_i(\succ_i, B_i))_{i \in I}$ be the resulting proxy preferences as defined above. Then we have $\varphi^{WP}(s) = \varphi^{GS}(P^*)$.

We next show that, while the USMA mechanism can be interpreted as an approximation to cosm, its outcome can be quite troubling even at equilibrium.

Example 4 (Equilibrium outcome of USMA mechanism can be unstable, unfair, and Pareto inferior to the outcome of cosm): There are 4 cadets i_1, i_2, i_3, i_4 and 4 branches b_1, b_2, b_3, b_4 with one slot each. Cadets are ordered as i_1, i_2, i_3, i_4 under the order-of-merit list π and $\lambda = 1$, so any cadet who has signed a branch-of-choice contract is given priority at each branch. Cadet preferences $(P_{i_1}, P_{i_2}, P_{i_3}, P_{i_4})$ are as follows:

$$\begin{aligned}
P_{i_1} &: (b_1, t_0) P_{i_1} (b_2, t_0) P_{i_1} (b_1, t_+) P_{i_1} (b_3, t_0) \cdots \\
P_{i_2} &: (b_1, t_0) P_{i_2} (b_1, t_+) P_{i_2} (b_2, t_0) P_{i_2} (b_2, t_+) P_{i_2} (b_4, t_0) P_{i_2} (b_4, t_+) P_{i_2} (b_3, t_0) P_{i_2} (b_3, t_+) \\
P_{i_3} &: (b_1, t_0) P_{i_3} (b_3, t_0) P_{i_3} (b_1, t_+) P_{i_3} (b_2, t_0) P_{i_3} (b_3, t_+) P_{i_3} (b_2, t_+) P_{i_3} (b_4, t_0) P_{i_3} (b_4, t_+) \\
P_{i_4} &: (b_4, t_0) P_{i_4} (b_2, t_0) P_{i_4} (b_4, t_+) P_{i_4} (b_2, t_+) \cdots
\end{aligned}$$

Consider the following strategy profile $(s_{i_1}, s_{i_2}, s_{i_3}, s_{i_4})$ under the USMA mechanism:

$$\begin{aligned}
s_{i_1} &= (\succ_{i_1}, B_{i_1}) && \text{with } b_1 \succ_{i_1} b_2 \succ_{i_1} b_3 \succ_{i_1} b_4 \text{ and } B_{i_1} = \{b_1\}, \\
s_{i_2} &= (\succ_{i_2}, B_{i_2}) && \text{with } b_1 \succ_{i_2} b_2 \succ_{i_2} b_4 \succ_{i_2} b_3 \text{ and } B_{i_2} = \{b_4\}, \\
s_{i_3} &= (\succ_{i_3}, B_{i_3}) && \text{with } b_1 \succ_{i_3} b_3 \succ_{i_3} b_2 \succ_{i_3} b_4 \text{ and } B_{i_3} = \{b_1\}, \\
s_{i_4} &= (\succ_{i_4}, B_{i_4}) && \text{with } b_4 \succ_{i_4} b_2 \succ_{i_4} b_1 \succ_{i_4} b_3 \text{ and } B_{i_4} = \{b_2\}.
\end{aligned}$$

The outcome of the USMA mechanism is:

$$\varphi^{WP}(s) = \{(i_1, b_1, t_+), (i_2, b_2, t_0), (i_3, b_3, t_0), (i_4, b_4, t_0)\}$$

Claim: Strategy profile $(s_{i_1}, s_{i_2}, s_{i_3}, s_{i_4})$ is a Nash equilibrium.

Proof of the Claim: Clearly cadet i_4 having received his top choice (b_4, t_0) cannot have a profitable deviation. Moreover, since the highest-priority cadet i_1 not only ranks branch b_1 as his first choice but also signs a branch-of-choice contract for branch b_1 under s_{i_1} , he will receive the only slot at branch b_1 no matter what other cadets do. Once branch b_1 is out of their reach, cadets i_2 and i_3 are each receiving their top choices among the remaining choices. Hence they cannot have a profitable deviation either.

Finally, consider cadet i_1 who receives his third choice (b_1, t_+) . Observe that he cannot receive his top choice (b_1, t_0) since cadet i_3 not only ranks branch b_1 as his first choice but also signs a branch-of-choice contract for branch b_1 under s_{i_3} . Cadet i_1 cannot receive his second choice (b_2, t_0) either. Suppose he did. Then since cadet i_4 ranks branch b_2 as his second choice while at the same time signing a branch-of-choice contract for b_2 under s_{i_4} , he would need to be assigned a seat at his top-choice branch b_4 . But with no slots left at branches b_2, b_4 , and having signed a branch-of-choice contract for b_4 under s_{i_2} , cadet i_2 would need to be assigned a slot at branch b_1 . That leaves only branch b_3 for cadet i_3 . But this is not possible, since cadet i_3 not only ranks branch b_1 as his first choice but also signs a branch-of-choice contract for it under s_{i_3} , leading to the desired contradiction. Hence, cadet i_1 does not have a profitable deviation either. \square

We next have several observations about the equilibrium outcome $\varphi^{WP}(s)$.

1. $\varphi^{WP}(s)$ is **not fair**. Even though cadet i_1 is the highest-priority cadet, he envies the assignment (b_2, t_0) of cadet i_2 . This situation happened because the USMA mechanism implicitly assumed that cadet i_1 prefers (b_1, t_+) to (b_2, t_0) simply because cadet i_1 signed a branch-of-choice contract for branch b_1 . What is more striking is that, cadet i_1 has not made a mistake in his choice of strategy. Since s is a Nash equilibrium outcome, there is simply nothing he can do.
2. $\varphi^{WP}(s)$ is **not stable**. Cadet i_1 , branch b_2 and contract (i_1, b_2, t_0) is such that

$$(b_2, t_0)P_{i_1} \underbrace{\varphi^{WP}(s; i_1)}_{=(b_1, t_+)} \text{ and } (i_1, b_2, t_0) \in C_{b_2} \left(\underbrace{\varphi^{WP}(s)}_{= \{(i_2, b_2, t_0)\}} \cup \{(i_1, b_2, t_0)\} \right)$$

3. The equilibrium outcome $\varphi^{WP}(s)$ is **inefficient** in a very strong sense. Not only it is Pareto dominated by the outcome of the cadet-optimal stable mechanism

$$\varphi^{GS}(P) = \{(i_1, b_2, t_0), (i_2, b_1, t_+), (i_3, b_3, t_0), (i_4, b_4, t_0)\},$$

the two outcomes raise the exact same additional man-years for the Army. So even when the Armys welfare is taken into consideration, the equilibrium outcome $\varphi^{WP}(s)$ is very implausible. \diamond

In our next example, we show that a cadet might suffer from an improvement in his order-of-merit list standing at equilibria of the USMA mechanism. This example is a slight modification of Example 4.

Example 5 (Equilibrium outcome of the USMA mechanism may not respect improvements): There are 4 cadets i_1, i_2, i_3, i_4 and 4 branches b_1, b_2, b_3, b_4 with one slot each. Cadets are ordered as i_1, i_2, i_3, i_4 under the order-of-merit list π_1 and $\lambda = 1$, so any cadet who has signed a branch-of-choice contract is given priority at each branch. Cadet preferences $(P_{i_1}, P_{i_2}, P_{i_3}, P_{i_4})$ are as follows:

$$\begin{aligned} P_{i_1} : & \quad (b_1, t_0) P_{i_1} (b_2, t_0) P_{i_1} (b_1, t_+) P_{i_1} (b_3, t_0) \cdots \\ P_{i_2} : & \quad (b_1, t_0) P_{i_2} (b_1, t_+) P_{i_2} (b_2, t_0) P_{i_2} (b_2, t_+) P_{i_2} (b_4, t_0) P_{i_2} (b_4, t_+) P_{i_2} (b_3, t_0) P_{i_2} (b_3, t_+) \\ P_{i_3} : & \quad (b_1, t_0) P_{i_3} (b_3, t_0) P_{i_3} (b_1, t_+) P_{i_3} (b_2, t_0) P_{i_3} (b_3, t_+) P_{i_3} (b_2, t_+) P_{i_3} (b_4, t_0) P_{i_3} (b_4, t_+) \\ P_{i_4} : & \quad (b_4, t_0) P_{i_4} (b_2, t_0) P_{i_4} (b_4, t_+) P_{i_4} (b_2, t_+) \cdots \end{aligned}$$

Consider the following strategy profile $(s_{i_1}, s_{i_2}, s_{i_3}, s_{i_4})$ under the USMA mechanism:

$$\begin{aligned} s_{i_1} &= (\succ_{i_1}, B_{i_1}) && \text{with } b_1 \succ_{i_1} b_2 \succ_{i_1} b_3 \succ_{i_1} b_4 \text{ and } B_{i_1} = \{b_1\}, \\ s_{i_2} &= (\succ_{i_2}, B_{i_2}) && \text{with } b_1 \succ_{i_2} b_2 \succ_{i_2} b_4 \succ_{i_2} b_3 \text{ and } B_{i_2} = \{b_1, b_4\}, \\ s_{i_3} &= (\succ_{i_3}, B_{i_3}) && \text{with } b_1 \succ_{i_3} b_3 \succ_{i_3} b_2 \succ_{i_3} b_4 \text{ and } B_{i_3} = \{b_1\}, \\ s_{i_4} &= (\succ_{i_4}, B_{i_4}) && \text{with } b_4 \succ_{i_4} b_2 \succ_{i_4} b_1 \succ_{i_4} b_3 \text{ and } B_{i_4} = \{b_2\}. \end{aligned}$$

Just as in the case of Example 4, the strategy profile s is a Nash equilibrium¹⁸ and the outcome

¹⁸The only difference with Example 4 is that cadet i_2 signs two branch-of-choice contracts for b_1, b_4 under s_{i_2} rather than only for b_1 .

of the USMA mechanism is

$$\varphi_{\pi_1}^{WP}(s) = \{(i_1, b_1, t_+), (i_2, b_2, t_0), (i_3, b_3, t_0), (i_4, b_4, t_0)\}.$$

Next consider a slightly different scenario where cadets are ordered as i_2, i_1, i_3, i_4 under a different order-of-merit list π_2 . It is easy to verify that the strategy profile s remains a Nash equilibrium under this change, while the outcome of the USMA mechanism changes:

$$\varphi_{\pi_2}^{WP}(s) = \{(i_1, b_2, t_0), (i_2, b_1, t_+), (i_3, b_3, t_0), (i_4, b_4, t_0)\}.$$

Hence, the assignment of cadet i_1 improves as a result of a worse order-of-merit ranking π_2 , showing that the USMA mechanism may penalize cadets as a result of improved order-of-merit list standing. \diamond

So why is it that the USMA mechanism can result in such undesired outcomes, even at equilibria, although it can be interpreted as an approximation of the cadet-optimal stable mechanism? After all, not only are both mechanisms extensions of the celebrated Gale-Shapley agent-optimal stable mechanism, but they both use precisely the same branch priorities. The reason is that whereas the cadet-optimal stable mechanism is a direct preference revelation mechanism that maintains all desired properties of its predecessor, the USMA mechanism is an indirect mechanism that arguably uses a somewhat simpler strategy space but as a result needs to rely on proxy-preferences. And these proxy-preferences might be a very rough approximation of true preferences, for several reasons:

1. Truthful preference revelation over branches is no longer necessarily optimal under the USMA mechanism.
2. Even when cadets are entirely truthful about their preferences over branches, they are often unable to declare their preferences over branch-price pairs since under the induced proxy-preference relation the pair (b, t_+) is assumed to be ranked right after (b, t_0) for any branch b for which a cadet signs a branch-of-choice contract.
3. In general, it is in cadets' best interests to sign branch-of-choice contracts for only some of the branches for which they are potentially willing to pay the increased cost t_+ . Hence, several acceptable pairs with increased cost will be considered unacceptable under the proxy preferences.

4. Under the proxy-preferences, it is assumed that the “expensive” pair (b, t_+) is always worse than the “cheap” pair (b, t_0) . While this is a safe assumption for a majority of cadets, there can be a few who might prefer otherwise, for various reasons.

We will say that a strategy $s_i = (\succ'_i, B_i)$ **suggests compromise** if there are two branches b, b' such that $b \succ'_i b'$ and yet $b' \in B_i$ even though $b \notin B_i$. Under such a strategy, a cadet signs a branch-of-choice contract for branch b' but not for b even though branch b is more preferred under his submitted preferences. One potential explanation for the use of a strategy that suggests compromise is that the cadet might be misrepresenting his preferences and branch b' might be preferred to branch b under his true preferences. The use of such a strategy can be also justified when the cadet is truthful about his preferences if he is not willing to “give up” the opportunity to receive his less favorite branch b' as well as any other branch that he ranks between b and b' at the cheaper cost t_0 . In either case, the use of a strategy that suggests compromise is a clear evidence of a cadet having difficulty representing his true preferences under the USMA mechanism. In 2010, of the 475 cadets who signed at least one branch-of-choice contract, 40 of them submitted strategies that suggest compromise.

In addition to difficulties that arise because of the difference between proxy-preferences and true-preferences, the outcome of the USMA mechanism is likely to suffer from further issues when an equilibrium outcome cannot be reached. For example, a cadet might “overpay” for his assigned branch, paying the increased cost of t_+ even though there are lower-priority cadets who receive the same branch at the base cost of t_0 . Such out-of-equilibrium outcomes decrease cadet satisfaction even though they may potentially increase the man-year gains by the Army. So far we have not compared the two mechanisms from the Army’s perspective. In the next section, we will argue that, while the adoption of the cadet-optimal stable mechanism may potentially reduce the number of man-years gained via the mechanism for a given λ , it is in all parties’ best interests to modify the man-years gained by adjusting this parameter under the cadet-optimal stable mechanism.

4.1 The Impact of Mechanism Choice on the Welfare of the Army

So far we have analyzed the cadet-branch matching problem almost entirely from the perspective of the cadets. Clearly, the social planner of this problem, namely the Army, also cares about the outcome of this problem. The objective function of the Army in this context might be quite

complex, but it is safe to assume that it likely depends on at least three factors:

Cadet welfare: Clearly the Army's welfare is paternalistic and it increases in cadet welfare. After all the Army has replaced its old mechanism as part of three incentive packages that it refers to as the *Officer Career Satisfaction Program*. The cadet-optimal stable mechanism is an unambiguous winner based on this factor. It increases cadet welfare not only via a more favorable outcome, but also via entirely removing the need for gaming the system and thus simplifying the decision process. Moreover, all coordination failures resulting in losses in efficiency, as well as suboptimal strategy choices resulting in unfair outcomes are entirely eliminated further increasing cadet satisfaction under the cadet-optimal stable mechanism. Simply put, the cadet-optimal stable mechanism fits perfectly with the spirit of Officer Career Satisfaction Program.

Respecting order-of-merit list standings to a large degree: Under the USMA mechanism, cadets can influence priorities for only the last 25 percent of slots at each branch. This is evidence of the Army's desire for the mechanism to respond highly to the order-of-merit list. Both mechanisms use the exact same branch priority structure, and hence based on this factor we have a tie between the two mechanisms.

Number of man-years gained via the mechanism: Since the proxy preferences constructed by the USMA mechanism is only a rough approximation of the true preferences, it is difficult to compare the outcomes of the two mechanisms based on this factor. However, for overdemanding branches, all of the last 25 percent slots will be assigned to cadets at the increased cost of t_+ . There may be cadets who overpay for their assignment at out of equilibrium outcomes of the USMA mechanism. Such cadet mistakes make up one reason why the USMA mechanism might result in more cadets paying the increased cost t_+ than under the cadet-optimal stable mechanism. In 2010, out of 1039 cadets, 215 were assigned their branch at the increased cost t_+ . Among them, only 25 cadets overpaid for their assignment. Assuming that the proxy preferences are truthful, there would have been 190 cadets who would be charged the increased cost t_+ under the cadet-optimal stable mechanism. We would like to emphasize that this is a worst-case scenario for the cosm since cadets are likely to sign significantly more branch-of-choice contracts under the cosm. Even if there is a slight reduction in the number of cadets receiving their assignments at the increased cost t_+ under cosm, the difference is due to cadets who overpay under the USMA mechanism. A much better way to increase the man-years gained by the Army is to slightly increase the parameter λ . For example, increasing λ to 0.3 is more

than enough to offset a potential loss of 25 cadets who otherwise overpay under the USMA mechanism.

Access to true cadet preferences is important for the Army for other reasons as well. While the enlisted ranks of the U.S. military exhibit a high level of demographic diversity, the leadership of the military has remained demographically homogeneous. In 2006, while 31 percent of the enlisted ranks of the military were African American or Hispanic, only about 16 percent of all officers were African American or Hispanic, and only 5 percent of all Generals were African American or Hispanic (Lim et al. 2009). This is cause for concern, since diversity is highly valued in the military. In a recent Rand Corporation report prepared for the Office of the Secretary of Defense, Lim et al. (2009) conclude that the relative scarcity of minorities in Combat Arms branches of the Army is a potential barrier to improving demographic diversity in the senior officer ranks. In 2006, 80 percent of all Generals were from combat arm branches. Using 2007 Army ROTC data, Lim et al. (2009) show that while 58 percent of white cadets' submitted first choices were in Combat Arms, only 31 percent of African American cadets' first choices were in Combat Arms.¹⁹ They also report that minorities tend to rank lower on order-of-merit score and conclude that these numbers may not truly reflect a lack of interest in minorities for Combat Arms. The following quote is borrowed from Lim et al. (2009):

In this exploratory study, we have demonstrated that it is critical for the Army to increase minority representation in key career fields to improve the racial and ethnic diversity of its top military officers. But we also contend that there is a strong need for a more in-depth analysis of the Army branching process. If, as our study suggests, minorities are indeed self-selecting into career fields with relatively limited promotion opportunities, why are they doing so? On the one hand, minority cadets could truly prefer different career fields than white cadets. In this case, policy should focus on ways to make combat career fields more appealing to minorities. On the other hand, minorities may not really prefer support career fields but rather may reason that they lack the OML ranking to get a more competitive

¹⁹In addition to the USMA, the U.S. Army accesses officers through two other programs: Reserve Officer Training Corps (ROTC) and Officer Candidate School (OCS). The USMA is the “flagship” institution of the U.S. Army, and a disproportionate number of its graduates go on to become Colonels and Generals. ROTC, on the other hand, is the “workhorse” of the U.S. Army, producing over twice as many officers as the USMA. As in the case of the USMA mechanism, ROTC mechanism is also vulnerable to preference manipulation.

career field (or they forecast a low probability of success in that career field). In this case, minority cadets might desire a Combat Arms career field but may opt for their most-preferred Combat Support or Combat Service Support career field thinking that they would never get a top Combat Arms assignment.

The vulnerability of the ROTC mechanism to preference manipulation thus has adversely affected the authors' ability to prescribe an adequate policy recommendation in this important analysis. While Lim et al. (2009) focus on the ROTC mechanism, their study is as relevant to the case of the USMA, reinforcing the value of a strategy-proof mechanism.

5 Extended Model with Multiple Nested Layers

So far our focus has been cadet-branch matching. However, the cadet-optimal stable mechanism is a perfect candidate for any real-life application where allocation is mainly based on priorities although these priorities can be improved for some of the slots by paying a premium. In this section, we refer to cadets as agents and to branches as objects, and allow a potentially different set of terms along with a potentially different priority order for each object.

For each object $b \in B$, there are q_b copies of the object and the priority order $\pi_b : I \rightarrow \{1, \dots, n\}$ represents claims of agents over object type b . As in the case of the cadet-branch matching problem, each agent consumes only one object, and for any two agents i, j with $\pi_b(i) < \pi_b(j)$ agent i has higher claims on a copy of object b than agent j , other things being equal. While having high priority for an object type makes it easier to secure a copy, it is not the only way to receive one. An object b can be obtained under $k(b)$ different terms. Let $T_b = \{t_b^1, \dots, t_b^{k(b)}\}$ denote the set of terms or costs for object b . For each object b we assume that t_b^1 is the default “cheapest” term, and terms get more “expensive” in increasing index. As a result, each agent's assignment is an object-cost pair (b, t) where $b \in B$ and $t \in T(b)$. For each object b , copies of the object are divided in $k(b)$ segments with q_b^r copies in segment $r = 1, \dots, k(b)$ such that $\sum_{r=1}^{k(b)} q_b^r = q_b$. For each object b ,

1. the priority ranking π_b determines the claims on the first q_b^1 copies that will be awarded at the cheapest cost t_1 ,
2. while the priority ranking π_b is still used as a tie-breaker, those who are willing to pay at least t_2 receive higher priority for the next q_b^2 copies although the cost of these units

never exceeds t_2 ,

⋮

- r. while the priority ranking π_b is still used as a tie-breaker, those who are willing to pay at least t_r receive higher priority for the next q_b^r copies although the cost of these units never exceeds t_r ,

⋮

- k(b). while the priority ranking π_b is still used as a tie-breaker, those who are willing to pay t_k receive higher priority for the last $q_b^{k(b)}$ copies.

Each agent i has strict preferences P_i over all object-cost pairs, and we assume that there are enough copies of the objects; hence $\sum_{b \in B} q_b \geq n$.²⁰ A contract is an agent-object-cost triple (i, b, t) such that $t \in T_b$. Let X denote the set of all contracts.

The only delicate part of this straightforward extension is an adequate specification of the **chosen set of object b** from a set of contracts $X' \subseteq X$. Given the above **nested multi-layer** priority structure, we construct the chosen set $C_b(X')$ and the rejected set $R_b(X') = X' \setminus C_b(X')$ as follows:

Phase 0: Remove all contracts that involve another object b' and add them all to the rejected set $R_b(X')$. Hence each contract that survives Phase 0 involves object b .

Phase 1: For the first q_b^1 potential elements of $C_b(X')$, simply pick the contracts with highest priority agents under π_b one at a time. When multiple contracts of the same agent are available in X' , choose the contract with the cheapest cost t_b^1 and reject the others, adding them to $R_b(X')$. Continue until either all contracts in X' are considered or q_b^1 elements of $C_b(X')$ are determined. If the former happens terminate the procedure, and if the latter happens proceed with Phase 2.

In general for $\ell = 2, \dots, k(b)$,

Phase ℓ : There are two sub-phases in each Phase $\ell > 1$. In the first sub-phase only contracts with costs t_b^ℓ or more are considered, and the choice is made is entirely based on the priority order π_b among such contracts. There is a second sub-phase only if extra copies are left, and

²⁰For most purposes we can make this assumption without loss of generality since one of these objects might be a uniformly worst (null) object with capacity larger than n .

in this second sub-phase the choice is lexicographically made first by the cost (with higher cost contracts favored) and then by the priority order π_b . We next give a more detailed description of these sub-phases.

Phase ℓ .1: Only consider contracts with costs t_b^ℓ or more. Among them pick the contracts with higher-priority agents under π_b for the next q_b^ℓ elements of $C_b(X')$. When multiple contracts are available for the same agent, choose the one with lowest cost (which has to be more than $t_b^{\ell-1}$ being considered in this phase) and reject the others. Continue until either all contracts with costs $t_b^\ell, \dots, t_b^{k(b)}$ are considered in X' or $\sum_{r=1}^\ell q_b^r$ elements of $C_b(X')$ are determined. If the former happens proceed with Phase ℓ .2, and if the latter happens proceed with Phase $\ell + 1$.

Phase ℓ .2: By construction, all remaining contracts have cost $t_b^{\ell-1}$ or less. In this sub-phase contracts with higher costs are picked before contracts with lower costs, and only then is the priority order π_b followed. When multiple contracts are available for the same agent, the highest-cost contract is chosen and the others are rejected. Continue until either all contracts in X' are considered or $\sum_{r=1}^\ell q_b^r$ elements of $C_b(X')$ are determined. If the former happens terminate the procedure, and if the latter happens proceed with Phase $\ell + 1$ (unless $\ell = k(b)$ in which case terminate the procedure).

Lemma 1*: Nested multi-layer priorities satisfy the law of aggregate demand for each object b .

Lemma 2*: Elements of X are unilateral substitutes for each branch b under nested multi-layer priorities.

Given lemmata 1* and 2*, our extended model preserves all key aspects of the cadet-branch matching problem. In particular, the *cumulative offer algorithm* gives a stable allocation that is weakly preferred to any other stable allocation by any agent, and the resulting *agent-optimal stable mechanism* is strategy-proof, fair, and respects improvements.²¹

A relatively straightforward application of this extended model is assignment of parking space. Consider a college campus where there are several parking lots. Each parking lot is a different object type and agents have preferences over these lots. The use of priority ranking is widespread in the assignment of parking space, as is price discrimination via permits of different

²¹The last property needs to be modified since there are potentially multiple priority orders. Balinski and Sönmez (1999) consider this more general version.

cost. Hence an agent-optimal stable mechanism is a natural candidate for this problem. Our next application, presented in the next subsection, is somewhat more involved.

5.1 Application to Centralized School Admissions

To motivate this application, we need to describe some features of college admissions in Turkey. Admission to colleges is centralized and includes private universities in the system. The priority of each student for each department (including private university departments) is determined by a central planner via a national centralized exam together with some other factors such as high school grades. All students submit their preferences over individual departments to the central planner, and as shown by Balinski and Sönmez (1999) the central planner assigns students to departments via the Gale and Shapley (1962) *college-optimal stable mechanism*.

Private universities have been recruiting a large fraction of the best students in Turkey since the first private university, Bilkent University, was launched in 1986. This is because they have been offering full fellowships to the best students. Private universities treat the slots reserved for fellowship students at a department as a separate unit for the purposes of centralized admissions. Hence, there are different codes for fellowship slots and non-fellowship slots at any department. Indeed, there are four different codes for each department of Sabanci University: one code each for full fellowship slots, half fellowship slots, quarter fellowship slots, and regular slots. Technically speaking, this aspect of Turkish college admissions makes it an application of matching with contracts rather than a standard Gale-Shapley two-sided matching. For example, a student may rank Economics at Sabanci University with full fellowship as his first choice, Economics at Bosphorus University (a leading public university) as his second choice, Economics at Sabanci University with half fellowship as his third choice, etc. This application of matching with contracts, however, can be treated as a standard Gale-Shapley two-sided matching model, since fellowship slots are treated as if they were a different department.

In contrast to the flexibility that attracts students of all backgrounds to private universities, students at public universities uniformly pay a minimal tuition fee for their undergraduate education. Essentially all students are admitted to public universities with tuition-waiver. This makes public universities highly dependent on public funding, which makes it difficult for them to compete with the top private universities. One recently debated policy suggestion has been to reserve some of the seats at public universities for students who are willing to pay a

much higher tuition fee (one that is closer to the actual cost of the education).²² If adopted, this policy change will also break the asymmetry between private universities and public universities in terms of the student pools they can target.

How could one adjust the centralized student assignment mechanism as a result of such a policy change? When a private university cannot fill some of its regular slots, it does not make sense to treat them as if they were fellowship slots, as this may not be in the best interests of a private university. But if a public university cannot fill some of the slots reserved for high-tuition students, these slots can and probably should be used for regular tuition-waiver students. Hence if public Turkish universities reserve some of the seats for high-tuition students, the student-optimal stable mechanism with the above-described chosen sets would be a very plausible alternative. Analogous to the implementation in the cadet-branch matching problem, a fraction of the slots will be reserved for students who are willing to pay an increased cost and any extra unassigned high-tuition slots will be “converted” to regular slots to avoid wasting public resources. Multiple tuition levels can be considered without compromising the existence of a plausible mechanism by using the chosen sets introduced for our extended model.

6 Conclusion

Market design is an exciting field that seeks to provide practical advice for various resource allocation problems. Auctions have been employed to allocate radio spectrum, electricity, and timber involving hundreds of billions of dollars worldwide (Milgrom 2004). More recent applications of market design include school choice, kidney exchange (Roth, Sönmez, and Ünver 2004, 2005a,b, 2007, and Ünver 2010), course allocation (Sönmez and Ünver 2010 and Budish and Cantillon 2011), and internet ad auctions (Edelman, Ostrovsky, and Schwarz 2007 and Varian 2007). In this paper, we presented a new market design application of cadet-branch matching that emerged as part of the Officer Career Satisfaction Program that was designed by a group of economists and officers at West Point’s Office of Economic and Manpower Analysis to boost career satisfaction and retention. The cadet-optimal stable mechanism that we propose will not only eliminate many of the issues currently present in the USMA mechanism and thereby increase cadet satisfaction, but also provide the Army with potentially valuable estimates of

²²While college admissions are decentralized in the U.S., reserving some of the seats for high-tuition students is in the same spirit with state universities admitting out-of-state students at a much higher tuition level.

the full potential of the branch-of-choice incentive program. With these estimates in hand, the Army will have the ability to accurately calibrate parameters of the mechanism in order to achieve desired objectives in terms of gained man-years.

Our paper also highlights the power of the matching with contracts approach in the context of priority-based indivisible goods allocation. Tools provided by Hatfield and Milgrom (2005) and Hatfield and Kojima (2010) are useful in the formulation of complicated priority structures (such as the one used by the USMA in cadet-branch matching) that cannot be formulated via the student assignment models of Balinski and Sönmez (1999) and Abdulkadiroğlu and Sönmez (2003).

References

- [1] Abdulkadiroğlu, A., P. A. Pathak and A. E. Roth (2005), The New York City High School Match, *American Economic Review, Papers and Proceedings*, 95, 364-367.
- [2] Abdulkadiroğlu, A., P. A. Pathak, A. E. Roth, and T. Sönmez (2005), The Boston Public School Match, *American Economic Review, Papers and Proceedings*, 95, 368-371.
- [3] Abdulkadiroğlu, A. and T. Sönmez (2003), School Choice: A Mechanism Design Approach, *American Economic Review*, 93, 729-747.
- [4] Balinski, M. and T. Sönmez (1999), A Tale of Two Mechanisms: Student Placement, *Journal of Economic Theory*, 84, 73-94.
- [5] Budish, E. and E. Cantillon (2011), The Multi-Unit Assignment Problem: Theory and Evidence from Course Allocation at Harvard, Chicago Booth working paper.
- [6] Calsamiglia C., G. Haeringer, and F. Klijn (2010), Constrained School Choice: An Experimental Study, *American Economic Review*, 100(4), 1860-1874.
- [7] Crawford, V. P. and E. M. Knoer (1981), Job Matching with Heterogeneous Firms and Workers, *Econometrica*, 49, 437-450
- [8] Echenique, F. (2011), Contracts vs. Salaries in Matching, *American Economic Review*, forthcoming.

- [9] Edelman, B., Ostrovsky, M. and Schwarz, M. (2007), Internet Advertising and the Generalized Second Price Auction: Selling Billions of Dollars Worth of Keywords, *American Economic Review*, 97, 242-259.
- [10] Gale, D. and L. Shapley (1962), College Admissions and the Stability of Marriage, *American Mathematical Monthly*, 69, 9-15.
- [11] Haeringer G. and F. Klijn (2009), Constrained School Choice, *Journal of Economic Theory*, 144(5), 1921-1947.
- [12] Hatfield, J.W. and F. Kojima (2008), Matching with Contracts: Comment, *American Economic Review*, 98, 1189- 1194.
- [13] Hatfield, J.W. and F. Kojima (2010), Substitutes and Stability for Matching with Contracts, *Journal of Economic Theory*, 145, 1704-1723.
- [14] Hatfield, J.W. and S.D. Kominers (2011), Contract Design and Stability in Matching Markets, Harvard University and Stanford University working paper.
- [15] Hatfield, J. W. and P. R. Milgrom (2005), Matching with Contracts, *American Economic Review*, 95, 913-935.
- [16] Kelso, A. S. and V. P. Crawford (1982), Job Matchings, Coalition Formation, and Gross Substitutes, *Econometrica*, 50, 1483-1504.
- [17] Kesten, O. (2010), School Choice with Consent, *Quarterly Journal of Economics*, 125(3), 1297-1348.
- [18] Kominers, S.D. (2011), On The Correspondence of Contracts to Salaries in (Many-to-Many) Matching, Harvard University working paper.
- [19] Lim N., J.P. Marquis, K.C. Hall, D. Schulker D., and X. Zhuo (2009), Officer Classification and the Future of Diversity Among Senior Military Leaders A Case Study of the Army ROTC, RAND corporation report.
- [20] Milgrom, M. (2004), Putting Auction theory to Work, Churchill Lectures in Economics. Cambridge. Cambridge University Press.

- [21] Ostrovsky, M. (2008), Stability in Supply Chain Networks, *American Economic Review*, 98(3), 897-923.
- [22] Pathak, P. A. and T. Sönmez (2008), Leveling the Playing Field: Sincere and Strategic Players in the Boston Mechanism, *American Economic Review*, 98(4), 1636-1652.
- [23] Pathak, P. A. and T. Sönmez (2011), School Admissions Reform in Chicago and England: Comparing Mechanisms by their Vulnerability to Manipulation, NBER Working Paper 16783.
- [24] Roth, A. E. and E. Peranson (1999), The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design, *American Economic Review*, 89, 4, 748-780
- [25] Roth, A. E. and M. Sotomayor (1990), Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis, Econometric Society Monograph Series, Cambridge University Press, 1990.
- [26] Roth, A. E., T. Sönmez, and M. U. Ünver (2004), Kidney Exchange, *Quarterly Journal of Economics*, 119, 457-488.
- [27] Roth, A. E., T. Sönmez, and M. U. Ünver (2005), Pairwise Kidney Exchange, *Journal of Economic Theory*, 125, 151-188.
- [28] Roth, A. E., T. Sönmez, and M. U. Ünver (2005), A Kidney Exchange Clearinghouse in New England, *American Economic Review Papers and Proceedings*, 95-2: 376-380, May 2005.
- [29] Roth, A. E., T. Sönmez, and M. U. Ünver (2007), Efficient Kidney Exchange: Coincidence of Wants in Markets with Compatibility-Based Preferences, *American Economic Review*, 97-3: 828-851.
- [30] Sönmez, T. and M. U. Ünver (2010), Course Bidding at Business Schools, *International Economic Review*, 51(1): 99-123.
- [31] Sönmez, T. and M. U. Ünver (2010), Matching, Allocation, and Exchange of Discrete Resources, J. Benhabib, A. Bisin, and M. Jackson (eds.) *Handbook of Social Economics*, Elsevier.

- [32] Switzer, B.T. (2011), A Tale of Two Mechanisms: US Army Cadet Branching, Masters Thesis, Pontificia Universidad Católica de Chile.
- [33] Ünver, U. (2010), Dynamic Kidney Exchange, *Review of Economic Studies*, 77 (1): 372-414.
- [34] Varian, H. (2007), Position Auctions, *International Journal of Industrial Organization*, 25, 1163-1178.
- [35] Wardynski, C., D.S. Lyle, and M.J. Colarusso (2010), Towards a U.S. Army Officer Corps Strategy For Success: Retaining Talent, Strategic Studies Institute.
- [36] Westkamp, A. (2010), Market structure and matching with contracts, *Journal of Economic Theory*, 145(5), 1724 - 1738.

Appendix: Proofs

Proof of Lemma 1: By construction of the USMA chosen set, all contracts of a given cadet can be rejected from a branch only when it reaches full capacity. Hence the size of the USMA chosen set can never shrink as the set of available contracts grows. \diamond

Proof of Lemma 2: Let $x = (i, b, t) \in X'$ be the only contract in X' that involves cadet i and suppose $x \notin C_b(X')$. Let $X'' \supset X'$. We have two cases to consider:

Case 1: $t = t_0$. Since $(i, b, t_0) \notin C_b(X')$, we have

$$\left| \left\{ j \in I : (j, b, t_0) \in X' \text{ with } \pi(j) < \pi(i) \right\} \right| \geq (1 - \lambda)q_b, \text{ and}$$

$$\left| \left\{ j \in I : (j, b, t_0) \in X' \text{ and } \pi(j) < \pi(i) \right\} \cup \left\{ j \in I : (j, b, t_+) \in X' \right\} \right| \geq q_b.$$

Therefore $X'' \supset X'$ implies

$$\left| \left\{ j \in I : (j, b, t_0) \in X'' \text{ with } \pi(j) < \pi(i) \right\} \right| \geq (1 - \lambda)q_b, \text{ and}$$

$$\left| \left\{ j \in I : (j, b, t_0) \in X'' \text{ and } \pi(j) < \pi(i) \right\} \cup \left\{ j \in I : (j, b, t_+) \in X'' \right\} \right| \geq q_b$$

as well. Hence $x \notin C_b(X'')$.

Case 2: $t = t_+$. Recall that $(i, b, t_0) \notin X'$ for this case by assumption.²³ Therefore $(i, b, t_+) \notin C_b(X')$ implies

$$\left| \left\{ j \in I : \{(j, b, t_0), (j, b, t_+)\} \cap X' \neq \emptyset \text{ and } \pi(j) < \pi(i) \right\} \right| \geq q_b$$

which together with $X'' \supset X'$ implies

$$\left| \left\{ j \in I : \{(j, b, t_0), (j, b, t_+)\} \cap X'' \neq \emptyset \text{ and } \pi(j) < \pi(i) \right\} \right| \geq q_b.$$

Hence $x \notin C_b(X'')$ for this case as well, completing the proof. \diamond

Proof of Proposition 1: The chosen set for any branch b satisfies the unilateral substitutes condition by Lemma 2. Hence Proposition 1 is a corollary of Theorem 5 in (Hatfield and Kojima 2010). \diamond

Proof of Proposition 2: Let X' be the outcome of the cadet-optimal stable mechanism. Fix a cadet i and let $x = (i, b, t) \in X'$. Suppose that cadet i prefers contract $x' = (i, b', t')$ to x . Observe that contract x' must have been offered to but rejected by the cumulative offer algorithm. Therefore if $t' = t_+$, then any cadet who is assigned a slot at branch b' has strictly higher priority than cadet i by construction of $C_{b'}$. If on the other hand $t' = t_0$, then any cadet j with $(j, b', t_0) \in X'$ has higher priority than cadet i . In either case $X'(j)P_i X'(i)$ implies cadet j has higher priority than cadet i , completing the proof. \diamond

Proof of Proposition 3: For any branch b , the chosen set C_b satisfies unilateral substitutes condition by Lemma 2 and the law of aggregate demand by Lemma 1. Hence Proposition 3 is a corollary of Theorem 7 in Hatfield and Kojima (2010). \diamond

Proof of Proposition 4: Fix a cadet i and let π_1 be an unambiguous improvement for cadet i over π_2 .

Scenario 1: First consider the outcome of cosm under priority order π_1 . Recall that by Remark 1, the order of cadets making offers has no impact on the outcome of the cumulative offer algorithm. Therefore, we can obtain the outcome of $\varphi_{\pi_1}^{GS}$ as follows: First, entirely ignore cadet i and run the cumulative offers algorithm until it stops. Let X' be the resulting set of contracts. At this point, cadet i makes an offer for his first-choice contract x^1 . His offer may cause a chain of rejections, which may eventually cause contract x^1 to be rejected as well. If

²³Proof of Lemma 2 does not work for the substitutes condition due to failure of this statement along with the inequality following the statement.

that happens, cadet i makes an offer for his second choice x^2 , which may cause another chain of rejections, and so on. Let this process terminate after cadet i makes an offer for his k^{th} choice contract x^k . There may still be a chain of rejections after this offer, but it does not reach cadet i again. Hence cadet i receives his k^{th} choice under $\varphi_{\pi_1}^{GS}$.

Scenario 2: Next consider the outcome of $\varphi_{\pi_2}^{GS}$, which can be obtained in a similar way: Initially entirely ignore cadet i and run the cumulative offers algorithm until it stops. Since the only difference between two scenarios is cadet i 's standing in the priority list, X' will again be the resulting set of contracts. Next cadet i makes an offer for his first-choice contract x^1 . Since π_1 is an unambiguous improvement for cadet i over π_2 , precisely the same sequence of rejections will take place until he makes an offer for his k^{th} choice contract x^k . Therefore cadet i cannot receive a better contract than his k^{th} choice under $\varphi_{\pi_2}^{GS}$ (although he can receive a worse contract if the rejection chain returns back to him). \diamond

Proof of Proposition 5: Fix a cadet-branch matching problem. Let $s = (\succ'_i, B_i)_{i \in I}$ be a Nash equilibrium strategy profile under the USMA mechanism and $\varphi^{WP}(s)$ be the resulting equilibrium outcome. Construct the resulting proxy preferences $P^* = (P_i(\succ'_i, B_i))_{i \in I}$ as defined above.

First consider any cadet i who is assigned the pair (b, t_+) under $\varphi^{WP}(s)$ and suppose that cadet i would have been assigned the pair (b, t_0) had he not signed a branch-of-choice contract for branch b . That is, cadet i would have been assigned the pair (b, t_0) under $\varphi^{WP}(s_{-i}, s'_i)$ where $s_i = (\succ'_i, B_i)$ and $s'_i = (\succ'_i, B_i \setminus \{b\})$. Therefore since s is a Nash equilibrium strategy, we must have $(b, t_+)P_i(b, t_0)$ for any such cadet, for otherwise cadet i would have a profitable deviation.

We will prove the theorem by showing that each step of the USMA algorithm under the equilibrium profile s can be interpreted as a sequence of offers made to the cumulative offer algorithm under the proxy preference profile P^* . Recall that the order of offers is irrelevant under the cumulative offer algorithm by Remark 1.

Consider Step 1 of the USMA algorithm, except let cadets make their proposals to their top choices one at a time rather than simultaneously. Interpret each of these offers to be made at the base cost of t_0 . As offers arrive, each branch b decides whether to hold the offer or reject it, following the USMA branch priorities: When an offer comes to branch b from cadet i there are three possibilities:

1. There are fewer than $(1-\lambda)q_b$ higher-priority cadets under priority list π that are currently on hold by branch b . In that case branch b holds the offer from cadet i .

2. There are at least $(1 - \lambda)q_b$ higher-priority cadets under priority list π that are currently on hold by branch b . However, for the remaining slots, there are fewer than λq_b higher-priority cadets under the adjusted priority ranking π_b^+ that are currently on hold by branch b . In this case we will treat the following two cases separately:

- (a) $b \notin B_i$ (i.e. cadet i has not signed a branch-of-choice contract under s_i): In this case branch b holds the offer from cadet i right away.
- (b) $b \in B_i$ (i.e. cadet i has signed a branch-of-choice contract under s_i): In this case branch b first rejects the offer of cadet i although cadet i makes a second subsequent offer to branch b right away with the increased cost t_+ which is put on hold by branch b .

3. If neither of the first two cases applies, branch b rejects the offer from cadet i for good.

Observe that the exact same rejection decisions would have been given under the cumulative offer algorithm for proxy preference profile P^* for this first sequence of offers. Repeat the above procedure with subsequent steps of the USMA algorithm until the algorithm terminates, essentially reinterpreting the USMA algorithm as a cumulative offer algorithm for a specific order of offers. \diamond

Proof of Lemma 1*: By construction of the chosen set, all contracts of a given agent can be rejected by an object only when it reaches full capacity. Hence the size of the chosen set can never shrink as the set of available contracts grows. \diamond

Proof Lemma 2*: Let $x = (i, b, t_b^u) \in X'$ with $u \in \{1, \dots, k(b)\}$ be the only contract in X' that involves cadet i and suppose $x \notin C_b(X')$. Let $X'' \supset X'$. Since $(i, b, t_b^u) \notin C_b(X')$, we must have

$$\left| \underbrace{\left\{ j \in I : \{(j, b, t_b^1), \dots, (j, b, t_b^u)\} \cap X' \neq \emptyset \text{ and } \pi(j) < \pi(i) \right\}}_{=\bar{I}(X')} \right| \geq \sum_{r=1}^u q_b^r.$$

Moreover, for $u < k(b)$ and for all $\ell \in \{u + 1, \dots, k(b)\}$, we must have

$$\left| \bar{I}(X') \cup \left\{ j \in I : \{(j, b, t_b^{u+1}), \dots, (j, b, t_b^\ell)\} \cap X' \neq \emptyset \right\} \right| \geq \sum_{r=1}^{\ell} q_b^r.$$

Therefore, since $X'' \supset X'$,

$$\left| \underbrace{\left\{ j \in I : \{(j, b, t_b^1), \dots, (j, b, t_b^u)\} \cap X'' \neq \emptyset \text{ and } \pi(j) < \pi(i) \right\}}_{=\bar{I}(X'')} \right| \geq \sum_{r=1}^u q_b^r$$

and for all $\ell \in \{u + 1, \dots, k(b)\}$,

$$\left| \bar{I}(X'') \cup \left\{ j \in I : \{(j, b, t_b^{u+1}), \dots, (j, b, t_b^\ell)\} \cap X'' \neq \emptyset \right\} \right| \geq \sum_{r=1}^{\ell} q_b^r$$

if $u < k(b)$. Hence $x \notin C_b(X'')$, completing the proof. \diamond