Cadet-Branch Matching

TAYFUN SÖNMEZ
Boston College

Prior to 2006, the United States Military Academy (USMA) matched cadets to military specialties (branches) using a single category ranking system to determine priority. Since 2006, priority for the last 25 percent of the slots at each branch has been given to cadets who sign a branch-of-choice contract committing to serve in the Army for three additional years. Building on theoretical work of Hatfield and Milgrom (2005) and Hatfield and Kojima (2010), Sönmez and Switzer (2013) show that the resulting new matching problem not only has practical importance but also it fills a gap in the market design literature. Even though the new branch priorities designed by the Department of the Army fail a substitutes condition, the cumulative offer algorithm of Hatfield-Milgrom gives a cadet-optimal stable outcome in this environment. The resulting mechanism restores a number of important properties to the current USMA mechanism including stability, strategy-proofness and fairness which not only increase cadet welfare consistent with OCSP goals but also provides the Army with very accurate estimates of the effect of a change in the parameters of the mechanism on number of man-year gains by the branch-of-choice incentive program. This new application also shows that matching with contracts model have great potential to prescribe solutions to real-life resource allocation problems beyond domains that satisfy the substitutes condition.

Categories and Subject Descriptors: J.4 [**Social and Behavioral Science**]: Economics

General Terms: Economics; Theory

Additional Key Words and Phrases: Market design

1. INTRODUCTION

The *matching with contracts* model (Hatfield and Milgrom 2005) is widely considered as one of the most important advances of the last two decades in matching theory. This powerful model embeds Gale and Shapley (1962) *two-sided matching* model and Crawford and Knoer (1981) - Kelso and Crawford (1982) *labor market model*, and it has given impetus to a flurry of theoretical research as well as new applications of market design such as *matching with regional caps* by Kojima and Kamada (2013), *school choice with soft-caps* by Hafalir, Yenmez and Yildirim (2013), *matching with slot-specific priorities* by Kominers and Sönmez (2013), and *cadet-branch matching* by Sönmez and Switzer (2013), Sönmez (2013). In this letter we focus on the USMA cadet-branch matching application by Sönmez and Switzer (2013).

2. THE MODEL

A **cadet-branch matching problem** consists of a finite set of cadets, denoted $I = \{i_1, i_2, \ldots, i_n\}$, a finite set of branches $B = \{b_1, b_2, \ldots, b_m\}$, a vector of branch capacities $q = (q_b)_{b \in B}$, a set of "terms" or "prices" $T = \{t_1, \ldots, t_k\} \in \mathbb{R}^k_+$, a list of cadet preferences $\succeq = (\succeq_i)_{i \in I}$ over $(B \times T) \cup \{\emptyset\}$, and a list of base priority rankings

Author's address: Department of Economics, Boston College, Chestnut Hill, MA 02467. Phone: (617) 552-3690. Fax: (617) 552-2308. E-mail: sonmezt@bc.edu.

$\pi = (\pi_b)_{b \in B}$.

Here the function $\pi_b : I \to \{1, \ldots, n\}$ represents the **base priority ranking** at each branch $b \in B$ and $\pi_b(i) < \pi_b(j)$ means that cadet i has higher claims to a slot at branch b than cadet j, other things being equal. Often the base priority ranking will be the same across all branches. In those cases we will refer to this uniform priority ranking as the **order-of-merit list** and denote it by π^{OML}. We assume that the terms of a match increase in its index. That is, $t_1 < \cdots < t_k$. A cadet who is assigned the pair (b, t) commits to serving in the military for at least t years.

Cadet preferences over branch-price pairs are strict, and **separable** in B:

$$\forall i \in I, \forall b, b' \in B, \forall t, t' \in T, \qquad (b,t) \succ_i (b',t) \iff (b,t') \succ_i (b',t').$$

That is, each cadet has well-defined preferences over branches alone independent of the price. We also assume that remaining unmatched is the last choice for each cadet. Let $\mathcal{P}$ and $\mathcal{Q}$ denote the set of all strict preferences over $(B \times T) \cup \{\emptyset\}$, and B respectively.

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 $X \equiv I \times B \times T$ denote the set of all contracts. An **allocation** $X' \subset X$ is a set of contracts such that each cadet appears in at most one contract and no branch appears in more contracts than its capacity. Let $\mathcal{X}$ denote the set of all allocations. For any allocation $X' \in \mathcal{X}$ and cadet $i \in I$, let X'_i denote the **assignment** of cadet i under X'. Here $X'_i = (b, t)$ if $(i, b, t) \in X'$, and $X'_i = \emptyset$ if cadet i has no contract in X'.

An allocation X' is **fair** if

$$\forall i, j \in I, \qquad \underbrace{X'_j}_{=(b,t)} \succ_i X'_i \Rightarrow \pi_b(j) < \pi_b(i).$$

That is, a higher-priority cadet can never envy the full *assignment* of a lower-priority cadet under a fair allocation.

A **mechanism** is a **strategy space** S_i for each cadet i along with an **outcome function** $\varphi : \prod_{i \in I} S_i \to \mathcal{X}$ that selects an allocation for each strategy vector $(s_1, s_2, \ldots, s_n) \in \prod_{i \in I} S_i$. A **direct mechanism** is a mechanism where the strategy space is simply the set of preferences over outcomes for each cadet i. A direct mechanism is **fair** if it always selects a fair allocation. A direct mechanism φ is **strategy-proof** if truthful preference revelation is dominant strategy.

Given two lists of base priority rankings π^1, π^2, we will say that π^1 is an **unambiguous improvement** for cadet i over π^2 if (1) the standing of cadet i is at least as good under π^1_b as π^2_b for any branch b, (2) the standing of cadet i strictly better under π^1_b than π^2_b for some branch b, and (3) the relative priority between all other cadets remain the same between π^1_b and π^2_b for any branch b. A direct mechanism **respects improvements** if a cadet never receives a strictly worse assignment as a result of an unambiguous improvement.

3. THE USMA MECHANISM

For the case of USMA, $\sum_{b \in B} \geq |I|$. That is, the total number of slots across all branches is no less than the number of cadets. All cadets receive their branch assignment through this mechanism and they are not allowed to declare any brach unacceptable. There are only two terms for USMA; $T = \{t_1, t_2\}$ with $t_1 < t_2$. We refer t_1 as the **base price**, t_2 as the **increased price**, and any contract with increased price t_2 as a **branch-of-choice contract**.

The (post-2006) **USMA mechanism** is not a direct mechanism and the space is $S_i = \mathcal{Q} \times 2^B$ for each cadet i under this mechanism. That is, each cadet is asked to choose (1) a ranking of branches alone, and (2) a number of branches (possibly none) under the USMA mechanism.

Fix a problem and let (Q_i, B_i) be the strategy choice of cadet i under the USMA mechanism. USMA interprets Q_i as the preferences of the cadet i over branches alone. The interpretation of the second element is somewhat more delicate. For each branch $b \in B_i$, cadet i indicates a willingness to pay the increased price t_2 in exchange for favorable treatment for the last 25 percent of slots in this branch. Cadet i will need to pay the increased price only if he receives one of the last 25 percent of the slots for which he is favored.

For each branch $b \in B_i$, we will say that cadet i **signs a branch-of-choice contract**. The following construction simplifies the description of the outcome function. Given a strategy profile $s = (Q_i, B_i)_{i \in I}$, uniquely construct the **adjusted priority ranking** $\pi_b^+ : I \to \{1, \ldots, n\}$ for each branch b as follows: For any pair of cadets $i, j \in I$,

$$\begin{array}{ll} \pi_b^+(i) < \pi_b^+(j) & \text{if } b \in B_i \text{ and } b \notin B_j, \\ \pi_b^+(i) > \pi_b^+(j) & \text{if } b \notin B_i \text{ and } b \in B_j, \text{ and} \\ \pi_b^+(i) < \pi_b^+(j) \iff \pi_b(i) < \pi_b(j) & \text{otherwise.} \end{array}$$

That is, the adjusted priority ranking π_b^+ is consistent with the base priority ranking π_b, unless one of the cadets signs a branch of choice contract for branch b whereas the other one does not. In that case the cadet who signs the branch-of-choice contract is favored under the adjusted priority ranking.

The algorithm that determines the outcome of the USMA mechanism is reminiscent of the celebrated agent-proposing deferred acceptance algorithm (Gale and Shapley 1962) with one important difference. While each branch b relies on the base priority ranking π_b to evaluate proposing cadets for the first 75 percent of the slots, it uses the adjusted priority ranking π_b^+ for the remaining slots. In other words cadets receive favorable treatment for the last quarter of slots for branches they have signed a branch-of-choice contract. We are ready to formally define the outcome function ψ^{WP} for the USMA mechanism.

Let $\lambda \in [0, 1]$. For a given strategy profile $(Q_i, B_i)_{i \in I}$, the USMA mechanism determines its outcome with the following **USMA algorithm**:

Step 1**:** Each cadet i "applies" to his top-choice under Q_i.

Each branch b *holds* the top $(1 - \lambda)q_b$ candidates based on the base priority ranking π_b. Among the *remaining applicants* it *holds* the top λq_b candidates based on the adjusted priority ranking π_b^+. 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 under Q_i.

Each branch b reviews the new applicants along with those held from Step (k-1), and *holds* the top $(1-\lambda)q_b$ based on the base priority ranking π_b. For the remaining slots, branch b considers all remaining applicants and *holds* the top λq_b of them based on the adjusted priority ranking π_b^+. 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 is charged the base price t_1,

(2) any cadet who is assigned one of the last λq_b slots is charged
   (a) the increased price t_2 if he has signed a branch-of-choice contract for branch b, and
   (b) the base price t_1 if he has not signed a branch-of-choice contract for branch b.

Let $\psi^{\text{WP}}(s)$ denote the outcome of USMA mechanism under $s = (\succ'_i, B_i)_{i \in I}$.

Observe that when $\lambda = 0$, the second part of the strategy space becomes redundant and the USMA algorithm reduces to agent-proposing deferred acceptance algorithm. Hence the USMA mechanism is equivalent to agent-optimal stable mechanism for this special case. Indeed, since all base priorities are identical for the case of USMA, the USMA mechanism further reduces to simple serial dictatorship induced by order-of-merit list π^{OML}. Both of these mechanisms are very well-behaved: Not only do they always result in a fair allocation, but they also respect unambiguous improvements and they are strategy-proof.

For the case of USMA, $\lambda = 0.25$, and for positive λ the analysis of the USMA mechanism is somewhat more delicate. That is because not only may truthful branch preference revelation be suboptimal under the USMA mechanism, but also the optimal choice of branch-of-choice contracts is a challenging task. So why is a mechanism that is clearly based on strategy-proof agent-optimal stable mechanism distorting incentives? The reason for this failure is somewhat subtle, but the key observation is the following. Under the USMA algorithm, each cadet is considered for one of the "expensive" slots at a branch immediately after its slots at base price, and more importantly before given an opportunity to propose to the cheaper slots of a less desirable branch. This may cause an issue if a cadet prefers the less preferred branch at base price to the more desirable brach at the increased price. Loosely speaking the USMA mechanism tries to infer cadet preferences over branch-price pairs from their submitted preferences over branches alone along with signed branch-of-choice contracts, and essentially it assumes that the price consideration is always secondary to the branch assignment for every cadet. To the extent that this assumption fails, the USMA mechanism distorts incentives. The following proposition summarizes the key shortcomings of the USMA mechanism.

THEOREM 1. *(Sönmez and Switzer 2013) Truth-telling may not be an optimal strategy under the USMA mechanism. Furthermore, a Nash equilibrium outcome of the USMA mechanism can be unfair, Pareto inferior to a fair allocation, and may penalize cadets for unambiguous improvements.*

4. MATCHING WITH CONTRACTS

Fortunately it is possible fix the deficiencies of the USMA mechanism. This requires relating cadet-branch matching to a recent model which has received a lot of attention.

4.1 Choice Sets

The cadet-branch matching problem can be modeled as 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. The original matching with contracts models is a two-sided matching model, and as such, each branch (hospitals in Hatfield and Milgrom 2005) has preferences over sets of agent-term pairs. These hospital preferences induce a **choice set** from each set of contracts, and it is this choice set (rather than hospital preferences) that relevant for cadet-branch matching.

In the present framework, branches are not agents and they do not have preferences. However, branches have priorities over cadet-price pairs, and these priorities can also be represented via choice sets. This is the sense in which the cadet-branch matching problem is a special case of matching with contracts.

In general, the choice set of branch b from a set of contacts X' depends on the policy on who has higher claims for slots in branch b. We can represent the *USMA priorities,* or any other priority structure by adequate construction of choice sets.

For a given priority structure for branch b, let $C_b(X')$ denote the set of contracts chosen from $X' \subseteq X$ whereas $R_b(X') \equiv X' \setminus C_b(X')$ denote the **rejected set** from $X' \subseteq X$. Given a set of contracts $X' \subseteq X$, we can formally define *USMA choice set* as follows.

USMA Choice Set

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

***Phase 1*:** For the first $0.75q_b$ potential elements of $C_b(X')$, choose the contracts with highest-OML cadets one at a time. When two contracts of the same cadet are available, choose the contract with the base price t_1 and reject the other one. Continue until either all contracts are considered or $0.75q_b$ 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 $0.25q_b$ potential elements of $C_b(X')$, give priority to contracts with increased price t_2. Hence in this phase only consider branch-of-choice contracts and among them include in $C_b(X')$ the contracts with highest-OML cadets. If any cadet covered in Phase 2.1 has two contracts in X' reject the contract with the base price t_1. 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, 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 price

t_1. 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.

4.2 Stable Cadet-Branch Matching

Since the seminal paper of Gale and Shapley (1962), the *stability* condition has been central to the analysis of two-sided matching markets. In the context of cadet-branch matching, an allocation X' is **stable** if (1) no cadet or branch is imposed an unacceptable contract, and (2) there exists no cadet i, branch b, and contract $x = (i, b, t) \in X \setminus X'$ such that $(b, t)\ P_i\ X'(i)$ and $x \in C_b\big(X' \cup \{x\}$.

Three properties of choice sets, or equivalently branch priorities in our context, have played an important role in the analysis of matching with contracts.

Priorities satisfy the **irrelevance of rejected contracts** (IRC) condition if

$$\forall Y \subset X, \forall z \in X \setminus Y \qquad z \notin C(Y \cup \{z\}) \implies C(Y) = C(Y \cup \{z\}).$$

That is, the removal of rejected contracts shall not affect the choice set under the IRC condition.

Priorities satisfy the **law of aggregate demand** (LAD) condition for branch b if

$$X' \subset X'' \Rightarrow |C_b(X')| \leq |C_b(X'')|$$

That is, the size of the chosen set never shrinks as the set of contracts grows under the LAD condition.

Lemma 1. *The USMA priorities satisfy both IRC and LAD conditions.*

Of the three conditions, the third one plays an especially important role in the two-sided matching literature. Priorities satisfy the **substitutes** condition for branch b if for all $X' \subset X'' \subseteq X$ we have $R_b(X') \subseteq R_b(X'')$. That is, under the substitutes condition any contract that is rejected from a set X' is also rejected from any set X'' that contains X'. Substitutes together with IRC imply that the set of stable allocations is non-empty (Hatfield and Milgrom 2005, Aygün and Sönmez 2013).

The substitutes condition along with IRC condition have been very "handy" in the analysis of matching with contracts: Fixed-point techniques in lattice theory has strong implications under these conditions. In particular these conditions together assure the existence of a stable allocation, and that the set of stable outcomes is a lattice. Hatfield and Milgrom (2005) build their highly influential model around these conditions. A recent paper by Echenique (2012) questions the value added of the matching with contracts model.

Theorem 2. *(Echenique 2012) The matching with contracts model can be embedded within the Kelso and Crawford (1982) labor market model under the substitutes condition.*

Kominers (2012) extends this isomorphism to a many-to-many matching provided that the two sides of the market can sign at most one contract.

The substitutes condition is key for both the above isomorphism to hold. Indeed Echenique (2012) emphasizes 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. One of the conditions offered in Hatfield and Kojima (2010) is the following:

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'. While the lattice structure of the set of stable allocations no longer persists under the unilateral substitutes condition, Hatfield and Kojima (2010) shows that a number of important results survives this weakening of the substitutes condition.

It turns out that, the unilateral substitutes condition plays a key role in cadet-branch matching:

LEMMA 2. *While USMA priorities do not satisfy the substitutes condition, they satisfy the unilateral substitutes condition.*

4.3 Cumulative Offer Algorithm and COSM

We refer the agent-optimal stable mechanism as cadet-optimal stable mechanism (COSM) in the present context. The strategy space of each cadet is $\mathcal{P}$ under the COSM, and hence it is a direct mechanism.

Fix branch priorities (and thus the choices sets). Given a preference profile $P \in \mathcal{P}$, the following cumulative offer algorithm (COA) (Hatfield and Milgrom 2005) can be used to find the outcome of COSM.

***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 $b \in B \setminus \{b(1)\}$.

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 $\bigcup_{b \in B} C_b(A_T)$.

Sönmez and Switzer (2013) build on the following result to fix the deficiencies of the USMA mechanism.

THEOREM 3. *(Hatfield and Kojima 2010): Suppose the priorities satisfy the unilateral substitutes condition along with IRC. Then the COA produces a stable allocation that is weakly preferred by any cadet to any stable allocation. Moreover, if priorities also satisfy the LAD condition then the induced COSM is strategy-proof.*

Remark 1. *The IRC condition is implicit in Hatfield and Kojima (2010) and it is not explicitly stated. The above interpretation, which is necessary in our framework with branch priorities, is due to Aygün and Sönmez (2012).*

Remark 2. *The Echenique (2012) embedding does not work in our framework. Hence, the cadet-branch matching problem is an application of matching with contracts that is beyond the scope of Kelso and Crawford (1982) labor market model.*

Let φ^{USMA} denote COSM induced by USMA priorities. This mechanism, which is a very modest deviation from the USMA mechanism, fixes all previously mentioned deficiencies of the USMA mechanism.

Theorem 4. *(Sönmez and Switzer 2013) The outcome of φ^{USMA} is stable under USMA priorities and it is weakly preferred by any cadet to any stable allocation. Moreover φ^{USMA} is strategy-proof, fair, and respects improvements.*

REFERENCES

Aygün, O. and Sönmez, T. 2012. Matching with contracts: The critical role of irrelevance of rejected contracts. Boston College working paper #805.

Aygün, O. and Sönmez, T. 2013. Matching with contracts: Comment. *American Economic Review* 103, 5, 2050–2051.

Crawford, V. P. and Knoer, E. M. 1981. Job matching with heterogeneous firms and workers. *Econometrica* 49, 437–450.

Echenique, F. 2012. Contracts vs. salaries in matching. *American Economic Review* 102, 594–601.

Gale, D. and Shapley, L. 1962. College admissions and the stability of marriage. *American Mathematical Monthly* 69, 9–15.

Hafalir, I. E., Yenmez, M. B., and Yildrim, M. A. 2013. Effective affirmative action in school choice. *Theoretical Economics* 8, 325-363.

Hatfield, J. W. and Kojima, F. 2010. Substitutes and stability for matching with contracts. *Journal of Economic Theory* 145, 1704–1723.

Hatfield, J. W. and Milgrom, P. R. 2005. Matching with contracts. *American Economic Review* 95, 913–935.

Kamada, Y. and Kojima, F. 2013. Improving efficiency in matching markets with regional caps: The case of the Japan Residency Matching Program. Stanford University working paper.

Kelso, A. S. and Crawford, V. P. 1982. Job matchings, coalition formation, and gross substitutes. *Econometrica* 50, 1483–1504.

Kominers, S. D. 2012. On the correspondence of contracts to salaries in (many-to-many) matching. *Games and Economic Behavior* 75, 2, 984–989.

Kominers, S. D. and Sönmez, T. 2012. Designing for diversity in matching. Boston College working paper #806.

Sönmez, T. 2013. Bidding for Army career specialties: Improving the ROTC branching mechanism. *Journal of Political Economy* 121, 1, 186–219.

Sönmez, T. and Switzer, T. B. 2013. Matching with (branch-of-choice) contracts at the United States Military Academy. *Econometrica* 81, 2, 451–488.