Designing Organizational Versus Public Markets*

Jonathan M.V. Davis
University of Chicago

B. Pablo Montagnes
Emory University

October 28, 2015

Abstract

Matching problems within an organization are distinct from traditional applications in public markets. Organizational assignment problems are constrained by individual rationality so an organization may select any assignment that is acceptable to its members. We show that there are no guarantees that assignment mechanisms that respect preferences will perform well from an organization's perspective. In some cases, an organization can better achieve its objectives by ignoring preferences and randomly choosing assignments, even when market participants have preferences aligned with the organization's, have outside options, and have private information about match qualities.

Keywords. Matching, Organizations

JEL Codes. C78, D2

*We thank John Hatfield, Fuhito Kojima, Alex Teytelboym, and especially Scott Kominers for helpful comments.

1. Introduction

In this paper, we show that matching problems within an organization[1] are distinct from traditional applications in public markets. Our main finding is that assignment mechanisms which are responsive to market participants' preferences do not always best serve the organization's objective, even in situations that seem favorable to a decentralized approach. In some cases, an organization can better achieve its objectives by ignoring preferences and randomly choosing assignments, even when market participants have preferences aligned with the organization's preferences, have outside options, and have private information about match qualities.

Many markets are composed of autonomous agents who are free to contract or match outside of a centralized mechanism. We refer to such markets as *public markets.* As a result, market designers almost always restrict attention to mechanisms which determine matches using agents' preferences, and often they restrict attention to mechanisms which yield pairwise stable matches[2] in the sense that no unmatched pair of agents would prefer to be matched together over their assigned match. In their seminal paper documenting the redesign of the National Residency Matching Program, Roth and Peranson (1999) explain: "Perhaps the most important and least controversial empirical finding about centralized matching algorithms is that they are most often successful if the matchings they produce are stable" (p. 752). Kamada and Kojima (2015) consider a situation where the government can prohibit blocking pairs but argue that "an assignment that completely ignores participants' preferences would be undesirable" (p. 77) and argue that stability is justified for normative reasons.

The success of market design in public markets has led to greater interest in the use of economic design to solve a diverse set of organizational problems. The International Monetary Fund uses the deferred acceptance algorithm (Gale and Shapley, 1962) to assign new economists to research teams[3] (Barron and Várdy, 2004), Teach For America uses a variant of the deferred acceptance algorithm to assign teachers to schools in some regions (Davis, 2015), and the United States Military Academy uses a cumulative offer mechanism to assign cadets to branches (Sonmez and Switzer, 2013; Sonmez, 2013). In all of these applications, matches are determined entirely by agents' preferences with no consideration of the optimal allocation from the organization's perspective.

Individuals choose to join an organization and may have to pay adjustment costs to switch to a new organization (for example, searching for and changing to a new job). As a result, organizations may have more flexibility in choosing an assignment mechanism than clearinghouses in public markets. An extreme case is the military which can court martial cadets for insubordination. In general, organizations do not have complete authority over their members and cannot prevent their members from exiting the organization, but they can forbid or sanction participants in internal two-sided matching markets from matching outside of the centralized assignment process. More concretely, organizational market design problems are constrained by individual rationality (IR) constraints. In many real world applications, organizations enforce an assignment without any input from the agents. For example, employees are often assigned to a supervisor or a group within a company as the company sees fit.[4]

The appeal of adopting the insights of market design within organizations has merit. Organizations may find it desirable to determine matches using a decentralized mechanism in order to bring some of the benefits of markets in to the organization. One might anticipate advantages to this approach when:

1. agents and organizations value similar match features;
2. agents possess more information than the organization; and
3. agents have outside options.

When agents and organizations value similar match features there appears to be little tension between the organization's goals and its members' preferences. Dispersed private information is a classic argument in favor of decentralization (Hayek, 1945), with the logic being that the decision maker who best understands the benefits and costs of a decision will be best able to optimally make the decision. A feature of decentralized mechanisms is their ability to aggregate this dispersed information. When agents have stronger outside options organizations must keep them happier in order to compete for their ongoing membership, perhaps by providing them with more autonomy over their matches. We show that none of the three conditions above, alone or in combination, is guaranteed to improve the performance of preference respecting mechanisms. In some cases, they make decentralized outcomes worse.

The rest of this paper is organized as follows. Section 2 presents a model of organizational market design problems. Section 3 presents our results. Section 4 concludes.

2. Model

We consider a variant of a two-sided assignment game (Shapley and Shubik, 1972; Roth and Sotomayor, 1992). Call one side of the market "teachers" and the other side "schools". Suppose there are M teachers and N schools to be matched in a market without prices or transfers. Each of the possible $M \times N$ matches is associated with an indivisible output:

$$\alpha_{ij} = f(i, j), \tag{1}$$

where $f(\cdot)$ is an arbitrary production function, i indexes teachers and j indexes schools.[5] This yields the following potential output matrix:

$$A = \begin{bmatrix} \alpha_{11} & ... & \alpha_{1N} \\ ... & & ... \\ \alpha_{M1} & ... & \alpha_{MN} \end{bmatrix}.$$

Teachers and schools get the following utility from each potential match, respectively:

$$u_{ij} = \alpha_{ij} + \mu_{ij} - \underline{\mathrm{u}}_i, \tag{2}$$

$$v_{ij} = \alpha_{ij} + \nu_{ij}, \tag{3}$$

where $\underline{\mathrm{u}}_i$ captures a teacher specific outside option. We assume that a match is acceptable to a teacher or school if $u_{ij} \geq 0$ or $v_{ij} \geq 0$, respectively. Teacher preferences can be written in matrix form as $U - \underline{\mathrm{U}}$, where U represents their preferences over matches and $\underline{\mathrm{U}}$ captures each teacher's outside option. Similarly, school preferences can be denoted V. Teachers and schools observe each component of their utility, including the actual output of the match, α_{ij}.

An *Organizational Market Design Problem* is a two-sided assignment game without prices or transfers with two distinguishing features. First, a risk neutral organization provides a technology that makes matches productive. We assume that matches are unproductive without organizational support. Second, the organization supporting the match is invested in the resulting assignment. Throughout this paper we assume the organization's objective is to assign teachers to schools in order to maximize total output:

$$\begin{aligned}
&\max_{x_{ij}} \sum_{j=1}^{N}\sum_{i=1}^{M} \alpha_{ij}x_{ij}, \\
\text{subject to } &\sum_{j=1}^{N} x_{ij} \le 1 \text{ for } i = 1,\ldots,M, \qquad (4)\\
&\sum_{i=1}^{M} x_{ij} \le 1 \text{ for } j = 1,\ldots,N, \\
&x_{ij} \in \{0,1\}, \qquad (5)\\
&u_{ij}x_{ij} \ge 0, \qquad (6)\\
&v_{ij}x_{ij} \ge 0 \text{ for } i = 1,\ldots,M, j = 1,\ldots,N.
\end{aligned}$$

In general, organizations may have different objective functions. For example, a pharmaceutical company trying to match scientists to labs may care about the maximum match quality rather than the sum of all match productivities since the best match is most likely to result in a scientific breakthrough.[6]

Organizations have imperfect information about match qualities. Prior to the match, the organization observes $\hat{A}^k_{ij}$, where for each i and j:

$$\hat{\alpha}_{ij} = \alpha_{ij} + \zeta_{ij}, \tag{7}$$

where ζ_{ij} is a random variable assumed to be independent of α_{ij}. Organizations and market participants have asymmetric information whenever $\text{Var}(\zeta_{ij}) \neq 0$.

The organization must decide whether to implement a match based on its noisy information about the match outputs, $\hat{A}_{ij}$ or to delegate the assignment decision to the market participants. For example, the organization could delegate assignments using the deferred acceptance algorithm (Gale and Shapley, 1962) or with other mechanisms that will not always yield stable matches[7], like Top Trading Cycles (Shapley and Scarf, 1974; Abdulkadiroglu and Sonmez, 1999) or a random serial dictatorship for one side of the market. We say that an assignment mechanism *respects agents' preferences* whenever two agents who most prefer each other are always matched by the mechanism. Respect for agents' preferences is a weaker notion than stability. All stable mechanisms respect agents' preferences since the pair who most prefer each other are a blocking pair if not matched together. However, other non-stable mechanisms, like top trading cycles, also respect agents' preferences.

In general, organizations can assign some matches and delegate others. The organization's problem amounts to selecting a matching algorithm and a constraint matrix, C:[8]

$$C = \begin{bmatrix} c_{11} & \dots & c_{1N} \\ \dots & & \dots \\ c_{M1} & \dots & c_{mn} \end{bmatrix},$$

where, for all i and j, c_{ij} is 1 if the organization allows the match and 0 otherwise. The constrained assignment problem is then delegated to market participants. Teacher and school preferences over the constrained set are given by $U \circ C - \underline{\mathrm{U}}$ and $V \circ C$, respectively, where $\circ$ denotes the Hadamard Product.[9]

3. Results

We begin our analysis by showing that in general an organization will not want to delegate the assignment decision to market participants using an assignment mechanism that respects agents' preferences when the organization has perfect information about the match productivity matrix, A, and teacher and school preferences. With perfect information, the organization's problem is a classical assignment problem as specified by Luenberger (2003). An optimal assignment exists and can be solved relatively easily using well-known solution methods, such as the Primal-Dual Transportation Algorithm. Therefore, with complete knowledge about A and participation constraints, the organization can do no better than to implement the optimal assignment without input from the teachers and schools. However, some organizations may prefer to give their members autonomy over matches if the resulting match is still optimal. Result 1 shows that in general assignment mechanisms that respect agents' preferences will not yield the optimal assignment.

Result 1. *There exists no assignment mechanism that respects agents' preferences that always selects an organization's optimal assignment.*

We demonstrate Result 1 using a counterexample.

Example 1. *Lankford et al. (2002) (LLW) show that minority and poor students are frequently taught by the least skilled teachers. This could be particularly troubling from a social welfare perspective if disadvantaged students, who likely have less support and fewer resources at home on average, benefit most from having a high quality teacher. LLW conclude that any salary differentials across urban and suburban schools are insufficient to compensate teachers for the difficulties associated with teaching in an urban school.*

In this example, we present a simple model of the scenario LLW describe. Suppose the market consists of one good and one bad teacher, one wealthy and one poor school, and a single organization supporting matches. Moreover, suppose teaching at the wealthy school is easier because wealthy students have more enriching home environments. Consequently, both the good and bad teacher are more productive at the wealthy school, but the bad teacher's comparative advantage at the wealthy school is greater since students at the wealthy school are generally successful. The outputs associated with the four possible matches are given by:

$$A = \begin{array}{lcc} & \textit{Wealthy School} & \textit{Poor School} \\ \textit{Good Teacher} & \alpha & 1 \\ \textit{Bad Teacher} & 1 & 0 \end{array},$$

where $1 < \alpha < 2$*. The organization's optimal assignment is:*

$$\{\textit{(Good Teacher, Poor School),(Bad Teacher, Wealthy School)}\},$$

which yields a total output of 2.

Suppose schools only value the productivity of the matches, but teacher preferences are given by:

$$U = \begin{array}{lcc} & \textit{Wealthy School} & \textit{Poor School} \\ \textit{Good Teacher} & \alpha & 1 \pm \zeta_{12} \\ \textit{Bad Teacher} & 1 & 0 \end{array},$$

where α *is the same as in* A *and* ζ_{12} *equals* 1 *half the time and* -1 *the other half of the time. As an example,* ζ_{12} *could represent how happy the good teacher was when she arrived at her interview with the poor school.*

If the organization delegates the decision to the market participants using a mechanism which respects agents' preferences, the mechanism will select the optimal assignment when $\zeta_{12} = 1$, since the Good Teacher and Poor School most prefer each other. This assignment yields a total output of 2. When $\zeta_{12} = -1$, the mechanism selects the following assignment:

$$\{\textit{(Good Teacher, Wealthy School),(Bad Teacher, Poor School)}\},$$

since the Good Teacher and Wealthy School most prefer each other. This assignment yields a total output of α. Therefore, the expected output of the delegated assignment is $\frac{1}{2}2 + \frac{1}{2}\alpha$ which is less than the output of the optimal assignment.

The key insight underlying Result 1 and Example 1 is that ordinal preferences are determined by teachers' and schools' absolute advantage in matches, rather than their comparative advantage. Consequently, a match that respects agents' preferences will only be optimal when absolute advantage and comparative advantage coincide. Choosing according to absolute instead of comparative advantage ignores the effect of a match on the productivity of other matches and therefore imposes a *displacement externality* on the organization and the other members of the organization. Importantly, a match must be supported by an organization in order to be productive, so an organization is free to choose any assignment that is acceptable to its members.

3.1 Aligned Preferences

In the introduction, we hypothesized that an organization may find it beneficial to delegate the assignment when market participants value the same match features as the organization, the organization has imperfect information, and market participants have outside options. Result 2 shows that the first of these conditions can actually make things worse for an organization's bottom line.

Result 2. *Assignment mechanisms which respect agents' preferences may yield worse outcomes for an organization as market participants' preferences become more aligned with an organization's objectives.*

We demonstrate Result 2 using another counterexample. Example 2 shows that assignment mechanisms that respect agents' preferences will not always select the organization's optimal assignment even in environments where the market participants' preferences are highly aligned with the organizations. We formalize what we mean by highly aligned preferences with the following definition.

Definition 1. *Teachers and schools are said to have preferences* ***aligned***[10] *with the organization if their preferences over matches are determined entirely by the productivity of the matches, or more precisely:*

$$u_{ij} = \alpha_{ij} - u_i,$$

$$v_{ij} = \alpha_{ij}.$$

Assuming that agents have this type of aligned preferences should be favorable to a decentralized approach. Without aligned preferences, one could easily construct examples where a decentralized approach does terribly because agents want to minimize the organization's productivity. Combined with our assumption that the organization would like to maximize the sum of all match productivities, assuming market participants have aligned preferences is equivalent to assuming that the organization would like to select the utilitarian optimal match for its members.

The next example demonstrates where assignments from preference respecting mechanisms can go wrong from the organization's perspective when preferences are aligned.

Example 2. *Consider again the setup of Example 1. Assume preferences are aligned, so teachers and schools only value the productivity of the matches. Also, assume the utility from all agents' outside options is 0. The organization's optimal assignment is still:*

$$\{\textit{(Good Teacher, Poor School),(Bad Teacher, Wealthy School)}\},$$

which yields a total output of 2. If the organization delegates the decision to the market participants using a mechanism which respects agents' preferences, the assignment will be:

$$\{\textit{(Good Teacher, Wealthy School),(Bad Teacher, Poor School)}\},$$

since the good teacher and wealthy school most prefer each other. This assignment yields a total output of α where $\alpha < 2$. By comparison, when the wealthy teacher also values how happy she is at her interview with the poor school as in Example 1, the expected output of the delegated assignment *is $\frac{1}{2}2+\frac{1}{2}\alpha$. Therefore, the expected output of a delegated mechanism is actually lower when teachers are assumed to value the exact same match features as the organization.*

Example 2 demonstrates that mechanisms that respect agents' preferences will not always yield better matches from an organization's perspective when market participants value something similar to the organization. In fact, these mechanisms may yield worse matches from the organization's perspective when preferences are aligned.

Ordinal preferences are determined by teachers' and schools' absolute advantage in matches. When preferences are aligned but absolute and comparative advantage are maximized by different matches, the match from a mechanism that respects agents' preferences cannot be optimal. However, if preferences are not aligned, as in Example 1, the noise in teacher preferences can be beneficial to the organization if it causes the perceived absolute advantage to coincide with comparative advantage.

3.2 Imperfect Information

Examples 1 and 2 were extreme because they made the unrealistic assumption that the organization has perfect information about match productivities. In this section, we consider another, perhaps more realistic, extreme case, by assuming the organization has no information about match qualities. The organization can either impose a randomly chosen assignment or capitalize on the private information of the market participants by delegating the match with an assignment mechanism that respects preferences at the cost of ruling out organizationally stable matches that may conflict with agents' preferences. Result 3 shows that preference respecting mechanisms do not always dominate random assignment.

Result 3. *Random assignment can dominate an assignment mechanism that respects agents' preferences from a risk-neutral organization's perspective even if preferences are aligned.*

We demonstrate Result 3 using a counterexample. Example 3 shows that delegated matches may not perform well even when the organization has no information about the matches and market participants value the same match features as the organization.

Example 3. *Consider the setup of Example 2. Assume preferences are aligned and that the utility from all agents' outside options is 0. Suppose the market consists of a good and bad teacher, a* *wealthy and poor school, and a single organization supporting matches. The true outputs associated with the four possible matches are identical to the outputs in Example 2.*

If the organization delegates the decision to the market participants using a mechanism that respects agents' preferences, the assignment is:

$$\{\textit{(Good Teacher, Wealthy School),(Bad Teacher, Poor School)}\},$$

which yields a total output of α. If instead, the organization randomly selects the assignment, the assignment will be:

$$\{\textit{(Good Teacher, Wealthy School),(Bad Teacher, Poor School)}\},$$

with a total output of α half the time and:

$$\{\textit{(Good Teacher, Poor School),(Bad Teacher, Wealthy School)}\},$$

with a total output of 2 the other half of the time. The expected match quality is therefore $1 + \frac{\alpha}{2}$ which is greater than α because $\alpha < 2$. The random assignment mechanism dominates the delegated mechanism.

Example 3 demonstrates that delegating the assignment may not be optimal even in situations with highly dispersed information and aligned preferences.

This result is driven by two features of Example 3. Like Example 2, absolute and comparative advantage are maximized by different matches, so the preference respecting match is not optimal. Since there are only two possible matches that do not leave either teacher or school unmatched, the preference respecting match is the worst possible match. As a result, the delegated assignment mechanism is dominated in expectation by any mechanism that does not always respect preferences, including random assignment.

3.3 Outside Options

Another argument for delegating assignments is individual rationality (IR) constraints. Intuition suggests that the benefit of giving market participants more input in assignments should be greater when organizations must compete for membership. Result 4 shows that this is not always the case. In some cases, delegated assignments perform worse as outside options improve.

Result 4. *Preference respecting mechanisms may perform worse when outside options improve.*

We demonstrate Result 4 using another counterexample. Unlike our previous examples, we construct Example 4 so that absolute and comparative advantage coincide when there are no viable outside options. However, when outside options improve, matches with relatively low absolute advantage become unacceptable. Only considering absolute rather than comparative advantage ignores the effect of a match on other match productivities. This imposes a displacement externality on the organization and induces the teacher with the lower match quality to leave the organization. In short, preference respecting mechanisms may cater too much to some members at the expense of other members.

Example 4. *Consider a similar setup to Example 2. Again, assume preferences are aligned and that the utility from all agents' outside options is 0. Suppose the market includes a good teacher and a bad teacher, a wealthy school and a poor school, and a single organization supporting matches. The outputs associated with the four possible matches are given by:*

$$A = \begin{array}{lcc} & \textit{Wealthy School} & \textit{Poor School} \\ \textit{Good Teacher} & 1.5 & 1 \\ \textit{Bad Teacher} & 1 & 0.75. \end{array}$$

If the organization delegates the decision to the market participants using a preference respecting mechanism, the unique assignment is:

$$\{\textit{(Good Teacher,Wealthy School),(Bad Teacher,Poor School)}\},$$

which yields a total output of 2.25. This is also the optimal assignment.

Now, suppose the utility from both teachers' outside options improves from 0 to 1. The bad teacher *no longer finds being matched to the poor school acceptable and will quit, so the assignment's output falls to 1.5. If instead, the organization randomly selects the assignment, the assignment will be:*

$$\{\textit{(Good Teacher, Wealthy School),(Bad Teacher, Poor School)}\},$$

with a total output of 1.5 half the time and:

$$\{\textit{(Good Teacher, Poor School),(Bad Teacher, Wealthy School)}\},$$

with a total output of 2 the other half of the time. The random assignment mechanism dominates the delegated mechanism with the improved outside options.

While Result 4 shows that better outside options are not necessarily favorable to a decentralized approach from an organization's perspective, this is in part due to our assumption of a fixed pool of potential members. To see the limitations of this assumption, consider the assignment of airline seat upgrades (Kominers and Sonmez, 2014). This could be considered a type of organizational market design problem since ticketed passengers on a flight have already purchased their tickets for the flight and face a large cost of switching airlines for that flight. Consequently, the airline is constrained by IR constraints when assigning seat upgrades for that flight and does not necessarily need to assign upgrades in accordance with passengers' preferences. However, the airline would be naive to ignore customer satisfaction because the passengers can easily switch to a new airline for their next flight. Moreover, future frequent travelers may avoid that airline if they know they are more likely to receive upgrades from a competing airline.

3.4 More General Solutions

The previous examples show that preference respecting assignment mechanisms do not always yield assignments that are desirable to an organization, even under conditions that seem quite favorable to a decentralized approach. More generally, organizations are not limited to a choice between choosing all assignments or delegating all assignments using a preference respecting mechanism. Instead, organizations can determine assignments using a decentralized, but constrained, solution, where matches are determined using a preference respecting mechanism, but market participants have constrained choice sets. The following example demonstrates how this approach may be optimal for an organization.

Example 5. *Suppose now that the market includes a single good teacher, two bad teachers, two wealthy schools, and one poor school. The true output matrix is given by:*

$$A = \begin{array}{cccc} & \textit{Wealthy School 1} & \textit{Wealthy School 2} & \textit{Poor School 3} \\ \textit{Good Teacher 1} & 12+\eta_{11} & 12+\eta_{12} & 9+\eta_{13} \\ \textit{Bad Teacher 2} & 9+\eta_{21} & 9+\eta_{22} & 3+\eta_{23} \\ \textit{Bad Teacher 3} & 9+\eta_{31} & 9+\eta_{32} & 3+\eta_{33} \end{array},$$

where the η_{ij} are independent and identically distributed with $P(\eta_{ij} = 0.5) = \frac{1}{2}$ and $P(\eta_{ij} = -0.5) = \frac{1}{2}$. The organization knows the mean of each match quality and the distribution of η_{ij}, but does not observe the η_{ij}. All expected outputs mentioned below are derived in Appendix 1.

If the organization delegates the assignment, the assignment will match the good teacher and one of the bad teachers with the wealthy schools and the remaining bad teacher with a poor school. In expectation, this will result in an aggregate output of 24.5. In contrast, the expected output of a randomly selected assignment is 25*. Again, the random assignment mechanism outperforms the delegated mechanism.*

However, the organization can do better than random assignment since it knows the expected productivity of each match. In particular, notice that the good teacher's comparative advantage is always greatest at the poor school.[11] *Since the optimal assignment will always place teachers in the position where their comparative advantage is greatest, the organization should optimally assign the good teacher to the poor school.*

After assigning the good teacher to the poor school, the organization must decide whether to randomly assign matches between the bad teachers and wealthy schools or to delegate these assignment decisions. The expected output from randomly choosing how bad teachers and wealthy schools are matched is 27, whereas the expected output of the delegated assignment is $27\frac{5}{16}$. The delegated assignment nearly achieves the optimal match's expected output of $27\frac{3}{8}$.

The organization's optimal solution to this Organizational Market Design Problem is to select the constraint matrix, C^*:

$$C^* = \begin{bmatrix} 0 & 0 & 1 \\ 1 & 1 & 0 \\ 1 & 1 & 0 \end{bmatrix},$$

and allow the market participants to select matches from their constrained choice sets using a delegated mechanism.

Our previous examples demonstrated that preference respecting assignments may be suboptimal because choosing assignments according to absolute advantage rather than comparative advantage imposes a displacement externality on the organization and the other members of the organization. In some cases, an organization may recognize that comparative advantage is maximized for a subset of its members by a particular match. In these cases, the organization can require the aforementioned subset of its members to match but allow its other members to match using a decentralized assignment mechanism. Example 5 demonstrates that while delegated assignments may perform worse than a completely random assignment in these cases, a constrained version of the delegated assignment may still be optimal.

4. Conclusion

Since the successful redesign of the National Residency Matching Program (Roth and Peranson, 1999) the tools of market design have increasingly been adopted as solutions to real world matching problems in situations without prices. Recent applications include the use of ordinal assignment mechanisms to solve matching problems within organizations (Barron and Várdy, 2004; Sonmez and Switzer, 2013; Sonmez, 2013). In this paper, we show that matching problems within an organization are distinct from traditional applications in public markets. In particular, organizations have flexibility in choosing an assignment mechanism that best achieves its objectives and may benefit from disregarding agents' preferences when picking assignments.

Our main finding is that assignments resulting from mechanisms that respect agents' preferences, like the deferred acceptance algorithm and top trading cycles, do not always best serve the organization's objective, even in situations that seem favorable to the decentralized approach. We illustrate examples where an organization is better off randomly choosing assignments even when market participants value the same types of things as the organization, have more information about match qualities than the organization, and have increasingly viable outside options.

References

Abdulkadiroglu, A. and Sonmez, T. (1999). House Allocation with Existing Tenants. *Journal of Economic Theory*, 88:233–260.

Barron, G. and Várdy, F. (2004). The Internal Job Market of the IMF's Economist Program.

Davis, J. M. (2015). Do Assignment Mechanisms Matter? Quasi-Experimental Evidence from a Teach for America Pilot.

Gale, D. and Shapley, L. S. . (1962). College Admissions and the Stability of Marriage. *The American Mathematical Monthly*, 69(1):9–15.

Hayek, F. (1945). The Use of Knowledge in Society. *The American Economic Review*, 35(4):519–530.

Jackson, C. K. (2013). Match Quality, Worker Productivity, and Worker Mobility: Direct Evidence from Teachers. *The Review of Economics and Statistics*, 95(4):1096–1116.

Kamada, Y. and Kojima, F. (2015). Efficient Matching Under Distributional Concerns: Theory and Applications. *American Economic Review*, 105(1):67–99.

Kojima, F. (2012). School choice : Impossibilities for affirmative action. *Games and Economic Behavior*, 75:685–693.

Kominers, S. D. and Sonmez, T. (2014). Matching with Slot-Specific Priorities : Theory.

Lankford, H., Loeb, S., and Wyckoff, J. (2002). Teacher Sorting and the Plight of Urban Schools : A Descriptive Analysis. *Educational Evaluation and Policy Analysis*, 24(1):37–62.

Luenberger, D. G. (2003). *Linear and Nonlinear Programming: Second Edition.* Springer, New York.

MacDonald, G. M. and Markusen, J. R. (1985). A Rehabilitation of Absolute Advantage. *Journal of Political Economy*, 93(2):277–297.

Niederle, M. and Yariv, L. (2009). Decentralized Matching with Aligned Preferences.

Roth, A. E. and Peranson, E. (1999). The Redesign of the Matching Market for American Physicians : Some Engineering Aspects of Economic Design. *The American Economic Review*, 89(4):748–780.

Roth, A. E., Rothblum, U. G., and Vande Vate, J. H. (1993). Stable Matchings, Optimal Assignments, and Linear Programming. *Mathematics of Operations Research*, 18(4):803–828.

Roth, A. E. and Sotomayor, M. (1992). Two-Sided Matching. In Aumann, R. and Hart, S., editors, *Handbook of Game Theory*, volume 1, chapter 16, pages 485–541. Elsevier Science Publishers B.V.

Shapley, L. and Scarf, H. (1974). On Cores and Indivisibility. *Journal of Mathematical Economics*, 1:23–37.

Shapley, L. S. and Shubik, M. (1972). The Assignment Game I: The Core. *International Journal of Game Theory*, 1:111–130.

Sonmez, T. (2013). Bidding for Army Career Specialties : Improving the ROTC Branching Mechanism. *Journal of Political Economy*, 121(1):186–219.

Sonmez, T. and Switzer, T. B. (2013). Matching With (Branch-of-Choice) Contracts at the United States Military Academy. *Econometrica*, 81(2):451–488.

Appendix 1. Example 5 Derivations

Expected Output from Unconstrained Mechanisms

In a completely random assignment, observe that each match is selected with probability $\frac{1}{3}$. Therefore, the expected output of a random assignment is $\frac{1}{3}(12 + 12 + 9 + 2(9 + 9 + 3)) = 25$.

If the assignment is delegated using a preference respecting mechanism, the good teacher will always match with one of the wealthy schools since $11.5 > 9.5$. This will be determined by the realizations of η_{11} and η_{12}. If η_{11} and η_{12} are both 0.5 or -0.5, then the good teacher will randomly select one of the schools and the output will be 12.5 or 11.5, respectively. If either η_{11} or η_{12} is 0.5 but the other is -0.5 the teacher will match with the more productive school. Therefore the expected output of the good teacher's match is $\frac{3}{4}12.5 + \frac{1}{4}11.5 = 12.25$.

One of the bad teachers will match with a wealthy school, say wealthy school 2, and the other will match with the poor school. Both bad teachers will prefer wealthy school 2 since $8.5 > 3.5$. If η_{22} and η_{32} are 0.5 or -0.5 wealthy school 2 will randomly pick one of the bad teachers and the output will be 9.5 or 8.5, respectively. If either η_{22} or η_{23} is 0.5 but the other is -0.5 the wealthy school will select the more productive match. Therefore the expected output of the bad teacher and wealthy school match is 9.25.

The remaining bad teacher, poor school match is not selected based on its productivity, so its expected output is 3.

Therefore, the expected output of the delegated assignment is 24.5.

Expected Output from Constrained Mechanisms

The good teacher is constrained to match with the poor school. The expected output of this match is 9.

If the bad teachers are randomly assigned to the wealthy schools, no information about the unobserved match components is taken in to consideration, so each match's expected output is also 9. The expected output of the constrained random assignment is 27.

Preferences will be determined by the realizations of η_{21}, η_{22}, η_{31}, and η_{32}. There are 16 possible realizations of the wealthy school-bad teacher matches. In 6 of these 16 cases, there is no variation in match quality across teachers or schools, so the output will be the same regardless of whether matches are determined by deferred acceptance, random assignment, or are optimally assigned. The expected output in these 6 cases is 18 since it is equally likely that both schools or both teachers have positive or negative draws.

If the realization is $(0.5, -0.5, -0.5, 0.5)$ or $(-0.5, 0.5, 0.5, -0.5)$, the bad teachers are well matched at different schools, so the delegated match will yield 19, as would the optimal match.

There are four cases where only one of the errors is positive. For example, $(0.5, -0.5, -0.5, -0.5)$. In these cases, the match with the positive draw will be most preferred by both the teacher and school, so it will always be included in the match resulting from the preference respecting mechanism. The positive and negative realizations will cancel out and the output from these matches will be 18. This is also the optimal assignment given these realizations.

In the final four cases, only one of the errors is negative. This is the only set of realizations where the optimal match may differ from the match resulting from the preference respecting mechanism. The optimal match will always include 2 positive draws and will have output 19. For our analysis of the mechanism, without loss of generality, assume that the realization is $(-0.5, 0.5, 0.5, 0.5)$. In this case, bad teacher 2 strictly prefers wealthy school 2 to wealthy school 1 and wealthy school 1 strictly prefers bad teacher 3 to bad teacher 2. Wealthy school 2 and bad teacher 3 are indifferent between teachers and schools, respectively. In the case where bad teacher 3 and wealthy school 2 choose each other, bad teacher 2 and wealthy school 1 will match and the output will be 18. In all other cases, bad teacher 2 will match with wealthy school 2 and the output will be 19. Therefore, the expected output of the delegated assignment is 18.75 $(\frac{1}{4}18 + \frac{3}{4}19)$.

Collecting the above cases, we see that the expected output of the constrained delegated assignment is $18\frac{5}{16}$ $(\frac{6}{16}18 + \frac{2}{16}19 + \frac{4}{16}18 + \frac{4}{16}18.75)$. The expected output of the optimal assignment is $18\frac{3}{8}$ $(\frac{6}{16}18 + \frac{2}{16}19 + \frac{4}{16}18 + \frac{4}{16}19)$.

Notes

[1]We refer to an "organization" instead of a "firm or organization" for ease of exposition.

[2]Following Kojima (2012) we refer to such mechanisms as "stable mechanisms".

[3]Barron and Várdy (2004) propose the use of deferred acceptance. We confirmed that it was implemented with Vardy on December 1, 2014.

[4]Roth et al. (1993) motivate their exploration of the relationship between matching and assignment problems with the example of matching workers to supervisors.

[5]Jackson (2013) studied teacher-school match qualities empirically by modeling $f()$ with a variant of the following Cobb-Douglas production function:

$$\alpha_{ij} = \gamma + \delta_i + \beta_j + \eta_{ij}, \tag{8}$$

where δ_i is a teacher-specific contribution to the match output, β_j is a school specific contribution to the match output, γ represents the organization's contribution the match output, and η_{ij} is a match specific component of productivity.

[6]MacDonald and Markusen (1985) present an example where optimal assignments within a firm are determined by the absolute advantage of workers. Since stability is determined by absolute advantage, stable mechanisms are likely to be optimal in their example.

[7]A match is stable if there is no unmatched school-teacher pair where both the school and the teacher would prefer to be matched together over their assigned match.

[8]There are $2^{N\times M}$ possible constraint matrices. Consequently, it becomes quite difficult to determine the optimal constraint matrix as the size of the market grows. In other work, we investigate the optimal choice of constraint matrix.

[9]The Hadamard Product multiplies matrices component by component.

[10]This is a special case of the definition used by Niederle and Yariv (2009) where the ordinal potential is given by the productivity matrix, A.

[11]The maximum comparative advantage the good teacher can have at the wealthy school is $12.5 - 8.5 = 4$, whereas the minimum comparative advantage she can have at the poor school is $8.5 - 3.5 = 5$.