

Redesigning the Israeli Psychology Master’s Match*

Avinatan Hassidim, Assaf Rommm, and Ran I. Shorrer†

January 16, 2017

Prior to 2014, the admission to Master’s and PhD programs in psychology in Israel was a mostly decentralized process. In 2013, in response to concerns about the existing procedure,¹ we proposed to design a centralized mechanism. Our goal was to apply the key lessons learned from decades of applied market design and use a mechanism that has good incentive properties. In particular, we wanted to use a mechanism that is both *stable* and *strategy-proof* for applicants. This paper describes how we successfully centralized this market, and the critical role of recent advances in matching theory.

During the design of the Israeli Psychology Master’s Match (IPMM), we met with the faculty of each of the participating programs and asked about the way they choose between applicants. We soon realized that their answers do not fit squarely into “traditional” models of two-sided matching markets. In particular, departments’ choice functions cannot be summarized by a quota and a rank-ordered list (ROL) for each program. Some departments employ affirmative action through minority quotas. Others aim to equalize the number of advisees each faculty member receives. And finally, some departments are willing to admit a limited number of applicants with different contractual terms (e.g., funding). Since terms can alter preferences between programs, this last feature implies that in order to satisfy the aforementioned desiderata, the applicants’ message space must be expressive enough to convey their preferences over program-terms pairs. This market is therefore a special case of the matching-with-contracts model (Hatfield and Milgrom, 2005).

Hatfield and Milgrom study an extension of Gale and Shapley’s (1962) college-admissions model, and show that a generalized version of the Deferred

* *Discussant*: Muriel Niederle, Stanford University.

† Hassidim: Bar-Ilan Univ., Ramat Gan, Israel, avinatanh@gmail.com. Romm: Hebrew Univ., Jerusalem, Israel, assaf.romm@mail.huji.ac.il. Shorrer: Penn State Univ., State College, PA 16801, USA. We thank Leon Deouell, Scott Kominers, Déborah Marciano, Al Roth, and Tayfun Sönmez for helpful discussions and suggestions. We thank the faculty and administrative staff of departments of psychology in Israel for their cooperation throughout the process. Hassidim is supported by ISF grant 1394/16. Romm is supported by a Falk Institute grant and ISF grant 1780/16.

¹The concerns largely mirror concerns about the decentralized matching process for American clinical psychologists in the 1990s (Roth and Xing, 1997).

Acceptance (DA) algorithm remains strategy-proof for applicants (Roth, 1982; Dubins and Freedman, 1981) and reaches a stable matching. These results rely on two conditions on colleges' choice functions: substitutability, meaning that options not chosen from a small menu will not be chosen from a larger menu, and the "law of aggregate demand" (LoAD), which assures that as the menu grows, the number of chosen options does not decrease.

Hatfield and Kojima (2010) found that the stability and strategy-proofness results go through if the substitutes condition is replaced by the weaker unilateral substitutes condition. Hatfield and Kominers (2015) later found an even weaker condition that ensures these properties. Up until recently, however, these findings had shed light on the structure of matching problems, but had not been applied in the field. Our experience designing the IPMM provides a strong empirical validation to the practical relevance of these recent results in matching theory.

The process of designing the IPMM took many months, during which we met with representatives from each department who explained their needs. In many cases, the described choice functions did not satisfy the substitutes condition, or even the unilateral substitutes condition. Yet, after each meeting we were glad to realize that there was a way to accommodate these choice function while assuring the good properties of DA. Only after the first year of operation did we learn about the findings of Hatfield and Kominers (2015), which explain why we were successful: the choice functions used by psychology departments have a substitutable completion that satisfies LoAD.² This ensures that DA converges to a stable matching and that it is strategy-proof.

The mechanism that we ended up implementing was indeed a variant of the applicant-proposing DA.³ Programs report their choice functions using a special interface that offers an expressive enough "bidding language" to report all of the pre-existing choice functions. Being able to accommodate departments' pre-existing choice functions was key for the implementation of our solution, as departments insisted on having the same flexibility they had under the previous process.

Our solution is similar to those proposed by Sönmez (2013) and Sönmez and Switzer (2013) for the allocation of cadets to military branches. In these two papers, the priorities to be used by the military are unilateral substitutes, but not substitutable. This is sufficient for finding more attractive alternatives to the mechanisms currently in use. To the best of our knowledge, the IPMM is the first field application of the Gale–Shapley program to a two-sided market with choice functions that violate even the unilateral substitutes condition. Therefore, this is the first documented market whose centralization required the full generality of the matching-with-contracts theory.

²In fact, during our investigation process one department discovered that its admissions criteria were not well defined. They then decided on a choice function that had a substitutable completion.

³Since 2015, application in couples has been supported through the Ashlagi, Braverman and Hassidim (2014) extension.

1 The psychology market prior to 2014

Admission to graduate degrees in psychology in Israel, especially in clinical programs, is highly competitive. The stakes are high: applicants have previously completed their Bachelor's studies in psychology, but this does not grant them the right to practice, and many of them seek a certifying clinical degree, which will eventually lead them to a career in a prestigious high-income profession. Each year there are about 1,400 new psychology graduates and fewer than 300 positions in clinical graduate programs. Other, non-clinical, graduate degrees are also available: about 300 applicants join these less demanded programs each year. Departments also compete over the best students.

Historically, the fierce competition induced a process of BA grades inflation and unraveling (i.e., increasingly earlier admissions). In an attempt to halt these processes, departments agreed to coordinate. They instituted a unified screening exam, set an earliest date for commencing the screening of applicants, and put in place a detailed protocol regarding the admission schedule. According to this protocol, departments were allowed to contact applicants on three different occasions (rounds) during a three-weeks period:

- On the first day of the first week, programs called applicants to notify them of their admission, wait-list status, or rejection. Applicants then had to inform programs within three days about the rejection of offers, or the tentative acceptance of a single offer.
- On the first day of the second week, programs called previously wait-listed applicants and notified them of their admission, rejection, or wait-list status. Applicants again had three days to respond, and were allowed to withdraw their previous acceptance, and to accept (irrevocably) at most one offer.
- On the third and final week, programs called applicants on their wait-list and offered admission. Applicants could no longer withdraw previous acceptances, and could only irrevocably accept incoming offers. Offers in this stage were often “exploding” (had to be accepted or rejected immediately or within a short period of time).

While this process was a major improvement relative to prior market conditions, it was problematic in several respects. The first and most acute problem was that it left much room for strategic behavior. Departments preferred to fill their capacity on the first and second rounds, so as to avoid the need to recruit on the third round. This fear motivated the costly collection of information about applicants' likelihood to accept offers. It also drove departments to offer admission to more students than they wanted to accept, in the expectation that some offers would be rejected. Applicants faced similar strategic problems. For example, since accepting an offer in the second round was irrevocable, applicants who got wait-listed by their preferred program in the beginning of the second round and received an offer from another program they liked less faced

the strategic choice between the “riskier” option of waiting and the “safer” acceptance of the less-preferred alternative.

The pre-existing process was also associated with excessive administrative costs, high levels of stress experienced by applicants, and much distrust among departments. Since programs had an incentive to act early, the general sentiment in many departments was that other departments were “cheating” by approaching candidates (explicitly or implicitly) before the prescribed dates, and sometimes “poaching” assigned candidates. Additionally, there was no way to make sure that applicants were following the rules, and not holding more than one offer at any given time.

2 Choice functions

In the process of designing the new system, we interviewed the officials in charge of admissions in each of the participating institutions. Our main question was how they choose whom to admit. It is important to underscore that answers were given verbally, and were not limited to a “bidding language.” In some cases, we were told that departments would prefer to choose applicants who ranked the program first or at least among the top k alternatives in their ROL. In these cases, we reminded the officials that this was not possible under the existing protocol, and explained that our solution would not accommodate this request either. In what follows, we describe the choice functions that were communicated to us (after the clarification above, when necessary).

All programs started by specifying different tracks (if applicable), by providing an ROL, and by setting the quota for the program and for each of the tracks separately.

1. Four departments’ programs had responsive choice functions. That is, admission to each program was determined solely by the provided quota and ROL. There were a few other departments that used such choice functions for a subset of their programs.
2. One program offered a number of seats in the “regular” track and a different number in the “honors” track, for which only highly ranked students were eligible. Unfilled seats in one track could not revert to seats in the other.
3. Some departments offered programs with priority seats for certain (mutually exclusive) categories of applicants.
4. One department offered several programs, some with priority seats and some with a direct PhD option.
5. Two departments had programs that offered a limited number of scholarships to a subset of applicants (with priority to get a scholarship being consistent with the regular ranking).

6. Three programs offered a limited number of funded seats and in addition (unfunded) priority seats to certain groups of applicants.
7. Three programs had a choice function that took into account the academic advisor the applicant would be assigned to. This choice function could be expressed by labeling applicants by their potential advisors and allocating priority seats to each advisor.
8. One department offered degrees in several programs, and in each program four different tracks: unfunded MA and a direct PhD with three different levels of funding. Each study program had a quota, and there was a restriction on the total number of scholarships.

Choice functions of types (1)–(7) have a substitutable completion that satisfies LoAD. (1) is responsive. (2)–(7) are either slot-specific or task-specific choice functions, see Kominers and Sönmez (2016) and Hatfield and Kominers (2015). Functions of type (8) involves simultaneous restrictions that may lead to the inexistence of a stable matching. Luckily, over-demand in this market assured that many quotas were certain to be filled, and this made the relevant department essentially indifferent between this choice function and another that had a substitutable completion that satisfied LoAD.

Some of these choice functions may fail to be unilateral substitutable. A simple example is that under choice functions of type (2) the admission of certain students to the “regular” track depends on whether contracts with stellar students are available in the “honors” track, or only in the “regular” track. For brevity, we do not discuss other examples, but we note that they were significantly more challenging to address.

3 The IPMM

Every year, since 2014, the matching process begins with an online registration phase.⁴ During this phase, departments report all the available contracts for applicants to rank,⁵ and applicants provide their national identification number and additional personal information.

After this stage is completed, and after programs conduct interviews, applicants are asked to rank contracts (i.e., program-terms pairs). For example, an applicant can rank a funded position in program *A* over a position in program *B* over a non-funded position in program *A*. Applicants submit their ROLs online using a simple drag-and-drop interface.⁶ If an applicant submits an ROL that includes a position in some program in which she did not rank all contracts, a pop-up alert appears. This design feature is meant to mitigate the

⁴The IPMM uses a website and software created by us especially for this purpose. The (Hebrew) website can be found at <http://www.psychologymatch.org>. The sources for the C++ program that collects departments’ choice functions and for the Python script that performs the match can be found on Romm’s homepage.

⁵Contracts specified a course of study, terminal degree (MA or PhD), and funding terms.

⁶Screenshots are available in an online appendix.

risk of applicants accidentally ranking only some of the positions offered by a program.⁷

Departments use a designated computer program to specify their choice functions. First, they can label applicants according to categories they choose to create (e.g., belonging to a minority, or being eligible for funding). Then, the department designates priority seats, and sets quotas and nested quotas. Except where priorities are specified, admission is based on a program-specific ROL. This language allows expressing all of the choice functions described above and assures that reported choice functions satisfy substitutable completability and LoAD.

Departments and applicants are informed that their preferences will not be revealed to other departments or applicants (other than in the form of aggregate statistics). The only exception is that the contact information (but not ROLs) of unmatched applicants are transferred to programs that failed to fill their capacity using the match, or had open positions due to “no-shows.”

In order to educate participants about the match, prior to the first year of operation, faculty and staff in participating departments attended presentations in which both DA and the fact that it was strategy-proof for the applicants were covered in great detail, in the hope that these individuals could provide good advice to applicants. It was also explained that for the programs, untruthful reporting could, in theory, be beneficial, but that gaining something from such a misrepresentation usually requires extensive knowledge of others’ behavior. Finally, applicants participating in the match were advised on multiple occasions to submit their true preferences, and were told that reporting false preferences could only hurt them as compared to telling the truth.⁸

4 Conclusion

While only three years of operation have passed, it appears that our efforts have been fruitful. The unraveling process that was previously in place has come to a halt and, as a by-product, trust between colleagues is gradually being restored.⁹ In addition, applicants are generally satisfied with the redesign. In a survey that was conducted following the 2015 match, satisfaction with the matching platform received an average score of 8.1/10. By contrast, satisfaction with the entire admissions process got an average score of 4.7/10.

Our experience also highlights the benefits of having a well-developed theory of matching markets. The market design literature has, naturally, given much attention to large markets. But much of the potentials gains from market

⁷However, see Hassidim, Romm and Shorrer (2016) for a discussion on how, despite this feature, some applicants decided not to rank funded positions.

⁸This advice is communicated in all emails and letters received from the automated matching system or from the departments themselves.

⁹Additionally, when we initiated the re-design process, one department chair mentioned that he would be happy to receive a class of the same quality “without all the fighting and the waste of so much time in the process.” In retrospect, he claimed that incoming classes are better than in previous years.

design lies in smaller scale markets. In 2013, in the absence of “off the shelf” solutions, we spent much time and effort designing the IPMM. Advances in matching theory, the basic science behind the design of entry-level labor markets, substantially reduce such costs and facilitate future design efforts.

References

- Ashlagi, Itai, Mark Braverman, and Avinatan Hassidim.** 2014. “Stability in large matching markets with complementarities.” *Operations Research*, 62(4): 713–732.
- Dubins, Lester E., and David A. Freedman.** 1981. “Machiavelli and the Gale–Shapley algorithm.” *American Mathematical Monthly*, 88(7): 485–494.
- Gale, David, and Lloyd S. Shapley.** 1962. “College admissions and the stability of marriage.” *American Mathematical Monthly*, 69(1): 9–15.
- Hassidim, Avinatan, Assaf Romm, and Ran I. Shorrer.** 2016. “Strategic’ Behavior in a Strategy-proof Environment.” Working paper.
- Hatfield, John W., and Fuhito Kojima.** 2010. “Substitutes and stability for matching with contracts.” *Journal of Economic Theory*, 145(5): 1704–1723.
- Hatfield, John W., and Paul R. Milgrom.** 2005. “Matching with contracts.” *American Economic Review*, 95(4): 913–935.
- Hatfield, John W., and Scott D. Kominers.** 2015. “Hidden Substitutes.” Proceedings of the 16th ACM Conference on Economics and Computation, 37.
- Kominers, Scott Duke, and Tayfun Sönmez.** 2016. “Matching with slot-specific priorities: Theory.” *Theoretical Economics*, 11(2): 683–710.
- Roth, Alvin E.** 1982. “The economics of matching: Stability and incentives.” *Mathematics of Operations Research*, 7(4): 617–628.
- Roth, Alvin E., and Xiaolin Xing.** 1997. “Turnaround time and bottlenecks in market clearing: Decentralized matching in the market for clinical psychologists.” *Journal of Political Economy*, 105(2): 284–329.
- Sönmez, Tayfun.** 2013. “Bidding for army career specialties: Improving the ROTC branching mechanism.” *Journal of Political Economy*, 121(1): 186–219.
- Sönmez, Tayfun, and Tobias B. Switzer.** 2013. “Matching with (branch-of-choice) contracts at the United States Military Academy.” *Econometrica*, 81(2): 451–488.