Matching for the Israeli "Mechinot" Gap Year: Handling Rich Diversity Requirements

(Extended Abstract)

Yannai A. Gonczarowski* Lior Kovalio† Noam Nisan‡ Assaf Romm§

April 21, 2018

Abstract

We describe our experience with designing and running a matching market for the Israeli "Mechinot" gap-year programs. The main conceptual challenge in the design of this market was the rich set of diversity considerations. This market was run for the first time in January 2018 and matched 1,607 candidates (out of a total of 2,580 candidates) to 35 different programs.

1 Background

Israeli youth typically graduate from high school at the approximate age of 18 and are then required to enlist in the army for about two and a half years (currently a bit less for women and a bit more for men). There are several institutions that offer high-school graduates a "gap year" before starting the army service, mostly focusing on some combination of educational and volunteering activities. It turns out that a significant number of youths, especially from the stronger socio-economic classes, are interested in taking such a gap year. In addition to the core educational and volunteering activities, these gap years typically also help them become more mature and independent, increase their self-confidence and their ability to get along with their peers, and build up their character, strengthening them for the challenging army service that lies ahead.[1] It seems that the army also sees a benefit from these gap years, and hence it allows the participants to postpone their mandatory army service until the end of the gap year.

This manuscript discusses one of the most significant types of institutions that offer such gap years, called in Hebrew "Mechinot Kedam Tseva'iyot" (מְכִינוֹת קְדַם־צְבָאִיּוֹת), in short "Mechinot"

*Einstein Institute of Mathematics, Rachel & Selim Benin School of Computer Science & Engineering, and Federmann Center for the Study of Rationality, The Hebrew University of Jerusalem, Israel; and Microsoft Research. *E-mail*: yannai@gonch.name.

†Rachel & Selim Benin School of Computer Science & Engineering, The Hebrew University of Jerusalem, Israel. *E-mail*: lior.kovalio@gmail.com.

‡Rachel & Selim Benin School of Computer Science & Engineering and Federmann Center for the Study of Rationality, The Hebrew University of Jerusalem, Israel; and Microsoft Research. *E-mail*: noam.nisan@gmail.com.

§Department of Economics and Federmann Center for the Study of Rationality, The Hebrew University of Jerusalem, Israel. *E-mail*: assaf.romm@mail.huji.ac.il.

[1] The public opinion of these programs in Israel seems to be that these are first and foremost opportunities to contribute to society by volunteering for a long period of time, in many cases with quite underprivileged populations, and that the personal gain in building character etc. are "fortunate consequences for the altruistic participants" rather than the main reason for attending these programs.

(מְכִינוֹת), and in English “Pre-Military Academies” (where “pre-military” is used in a strictly chronological sense), abbreviated PMAs in this paper. The PMAs, in general, focus more time on study than on volunteering, with an emphasis on various issues related to Jewish thought and to Israeli society, where the study is “for the sake of studying”: with no grades and no certificates. There are over 50 different PMAs, and these are very heterogeneous: some are religious, some secular, and some mixed; some are co-ed and some (mostly the religious ones) are single-gender; they have different mixes of activities, different focuses of studies, different philosophical and social approaches (from purely religious orthodox to very progressive), and, one may comfortably say, quite different political leanings. Each PMA is independent and is separately run and administered, but they voluntarily cooperate with each other through “The Joint Council of Pre-Military Academies” (a cooperation that can be considered quite remarkable considering the extreme range of social and political leanings represented).[2]

Let us describe how the admissions process to the PMAs worked before 2018. During the Fall of their senior year of high school, candidates start considering which PMAs they may want to go to in the following year (as well as several other options, including other types of gap years and not taking a gap year at all). The most critical part of the admissions process is an actual visit of the candidate to the PMA, usually spending a weekend there, participating in activities, getting to know the place, and getting to be known by the PMA staff. These visits go on more or less continuously during the months from October to January.

The PMAs see this inflow of candidates and *build a group* for the upcoming year: each PMA has a target number of candidates that it wishes to accept for the upcoming year — usually between 25 and 100 — which is determined by various constraints, such as physical space and educational agenda and, perhaps most significantly, by the number of “army-service deferments” granted by the army (since each candidate must defer his or her service by a year). The considerations employed by the PMAs are rather complicated and mirror their educational mission. On the individual level the PMAs naturally desire the “standard qualities” such as willingness to learn and volunteer, but they also value affinity to their own educational agenda and special character. More interestingly, most PMAs view their educational agenda also on a societal level, and so see being a meeting place for of the wide spectrum of the Israeli society as a central part of their mission. The PMAs then build their groups incrementally throughout from October throughout January, accepting candidates that they desire as these candidates visit the PMA, at each point in time taking into account not only the candidate himself or herself, but also how he or she “fits into the group built so far.” This allows each PMA in an online fashion to correct any current gender imbalance, make sure not to create a concentration of too many candidates from a single high school or small town, implicitly maintain a wide variety of balances according to different social criteria that the PMA cares about, as well as various affirmative-action types of policies for various less or more well-defined fragments of the Israeli society.

So, during the several months of this admissions process (as it ran before our involvement), what transpired was a distributed online process where candidates continuously visit PMAs, and then the PMAs continuously accept or reject candidates. This is a difficult distributed task and indeed it had not been working well: since the PMAs build their group in an online fashion, they need to know which of the candidates that they accepted so far actually intends to come. This has led the PMAs a give “exploding offers” to the candidates, where a candidate often has to accept an offer within days or lose his or her place. The candidates find it very difficult to reply to these exploding offers as these come before they have even visited many of the other PMAs (indeed, it was not uncommon for an attractive candidate to receive an such an exploding offer from the first

[2] See `https://mechinot.org.il/en-us/`.

PMA that they visited) so they do not yet even have a clear picture of where they would like to go. These two conflicting constraints naturally led to much pressure, requests for extensions, strategic timing decisions, and ex-post change of heart on all sides. This process has lead to significant sub-optimality of the outcome both for the candidates who may need to reject an offer due to having previously accepted an inferior one, and for the PMAs who need to make acceptance decisions before they have even seen all of the candidates. If one adds to this the psychological burden since even a successful candidate will naturally get many rejections and be under significant uncertainty for a very long time, and additionally the unfortunate fact that there currently is a significant shortage of slots and so a large number of candidates will not get accepted to any PMA, one can understand the general dissatisfaction with this original system.

2 Our System

During the year 2017, the authors of this manuscript have approached the Joint Council of PMAs and suggested switching to a centralized computerized admissions system. As it turned out that the PMAs were well aware of the difficulties with the existing system, and they were rather happy to do so. However, it was critical that each PMA maintain full educational sovereignty as well as maintain an individual relationship with each candidate that they may accept. So the admissions process for the 2018 PMAs (starting in the fall) was done using a computerized system that was designed, built, and operated by the authors. This centralized matching excluded the religious PMAs (that have somewhat different circumstances) as well as a few small PMAs that cater to special segments of the population. Some of the PMAs offer more than a single "program" differing in character or in geographic location, and these different programs were viewed as separate PMAs by the matching system. Altogether the system handled 35 different such programs (administered by 24 independent PMAs), with a total of 1,760 slots.

The process worked as follows: from October 2017 through January 2018, the candidates visited the PMAs and "interviewed" there as before in a distributed online manner. However, the PMAs did not accept or reject any candidate during this period, but rather only noted an evaluation of the candidate as well as any special attributes that they may care about when "building a group" (e.g., city of residence, gender, being religious, etc.) Similarly, the candidates did not need to accept or reject any PMA, but simply noted to themselves how they evaluate each PMA that they visited. During the first half of January 2018, each candidate had to login into a system that we deployed and enter his or her ranking of the visited PMAs. Similarly, by mid January, every PMA had to enter its preferences into the system using a specially formatted Excel spreadsheet, which we describe below. Figure 1 shows a screenshot of the candidates' graphical user interface that allows them to create a ranking of PMAs by dragging and dropping their icons into an ordered list.[3] Figure 2 shows an example of a spreadsheet by which PMAs describe their preferences. Once all the candidate and PMAs preferences were in the system, a variant of the Deferred Acceptance (DA) algorithm of Gale and Shapley (1962) was used to compute an assignment of candidates to PMAs.

The main new ingredient that our system had to handle was the large set of "diversity" constraints that guide the PMAs when "building a group." There is a significant literature that attempts to deal with various types of "quotas" and affirmative-action-type constraints in Gale-Shapley-like matching markets.[4] A survey of many of the approaches for handling such issues can

[3]The candidates' graphical user interface was adapted from a similar system used for the Israeli Psychology Master's Match (Hassidim et al., 2017b).

[4]See, for example, Biró et al. (2010) and Fragiadakis et al. (2016) on minimum quotas, and Hafalir et al. (2013) and Kojima (2012) on affirmative action.

Figure 1: The graphical user interface by which a candidates specifies his or her preferences

be found in a recent paper by Nguyen and Vohra (2017). It seems, however, that the richness of the balance requirements that the PMAs desire does go significantly beyond previous analysis. Our approach is the following: we devise a “language” that allows the PMAs to express their concerns when choosing among a group of candidates. In our system, all quota, balance, diversity, and affirmative-action constraints are expressed within the preferences of each of the PMAs separately (in the terminology of Nguyen and Vohra (2017) this is the “modifying priorities” approach). As is common for such “bidding languages” (see, e.g., Nisan, 2006), the language must strike a compromise between different concerns: being expressive enough as to handle the real preferences of the PMAs on the one hand, and being simple enough to be handled well on the other hand. This simplicity in our case has a triple meaning: algorithmic simplicity of running the match; strategic simplicity of handling the incentives induced; and cognitive simplicity so that the humans running the PMAs can actually use it well.

Our language, encoded in the excel spreadsheet format used by the PMAs depicted in Figure 2, is the following: each PMA defines a set of “categories” that it cares about (e.g., “gender,” “religiousness,” “city”). For each category the PMA is allowed to define a *maximum quota* as well as *minimum target*. In addition to ranking the candidates individually, the set of categories to which

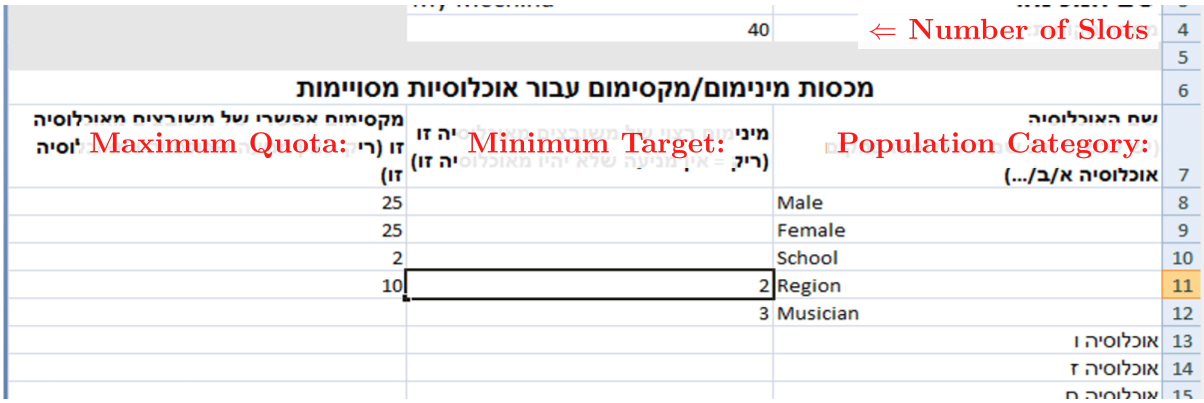

(a) Category input sheet

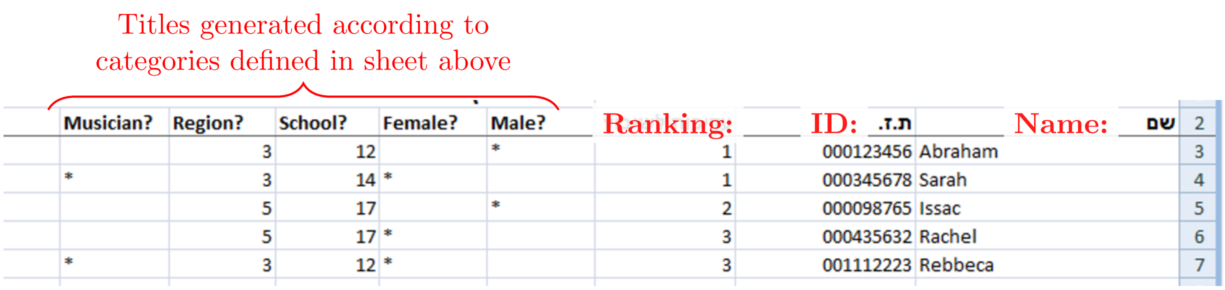

(b) Candidate input sheet

Figure 2: A spreadsheet by which a PMA describes its preferences (English translation overlayed)

each candidate belongs is specified.[5] Now, the interpretation of such a preference is the following: choose the top candidates according to their individual ranking, subject to the hard constraint that the set of candidates from any given category never exceeds the category's maximum quota, and strictly preferring (in a way that overrides individual ranking) candidates that belong to any category whose minimum target has not been reached. There are several issues that need be specified before this becomes fully formal, and these are specified in Algorithm 1.[6] It is worthwhile to mention though that in terms of cognitive simplicity, none of the PMAs showed any desire to dig into these issues beyond this (non-algorithmic) general verbal interpretation.

[5]As can be seen in Figure 2b, categories can be either "binary" (e.g., "Male" or "Musician") or "multiple-valued" (e.g., "School" or "Region"). For the latter, the minimum target and maximum quota are applied to each "value" (e.g., specific school or region) separately. So, a multiple-valued category with, say, 10 possible values, is no more than a shorthand for 10 disjoint binary categories. The ability to have multiple-valued categories is an example of a feature which, despite not extending the expressiveness of the bidding language from a computer-science perspective, makes the language much easier to use for the humans who input the data (and is less error-prone in terms of, e.g., accidentally skipping a line), and in fact is a feature that we added following the request of several PMAs while the system was already operational.

[6]Of note is our design decision to give the same "promotion" in the preferences of a PMA to candidates that belong to any (positive) number of categories whose minimum target has not been reached, regardless of the (positive) number of such categories to which such a candidate belongs. This decision was guided by the desire to prevent a *systematic* phenomenon where each of only a handful of candidates "fills up" many affirmative action slots, which could have resulted in a matching system biased toward having as few candidates as possible benefit from such slots.

```
input  : Set C of candidates to choose from
output: Set Ĉ of candidates chosen from C
// Initialization
Ĉ ← ∅;
// First pass: give higher priority to candidates who help reach minimum
   targets
foreach candidate c ∈ C, from highest-ranked to lowest-ranked do
    if c is in some category whose minimum quota is not met by Ĉ then
        if adding c to Ĉ does not violate any maximum quota then
            Ĉ ← Ĉ ∪ {c};
        end
    end
end
// Second pass: all other candidates
foreach candidate c ∈ C \ Ĉ, from highest-ranked to lowest-ranked do
    if adding c to Ĉ does not violate any maximum quota then
        Ĉ ← Ĉ ∪ {c};
    end
end
```

Algorithm 1: The *choice function* of a PMA, specifying for each set C of candidates, which subset $\hat{C} \subseteq C$ of these candidates were to be chosen by the PMA (according to the preferences declared in its Excel spreadsheet) if that PMA had free choice from C (and only from C). When traversing over candidates from highest-ranked to lowest-ranked, ties are broken randomly but consistently across schools (i.e., single tie-breaking, see Abdulkadiroğlu et al., 2009).

2.1 Algorithm and Theoretical Barriers

As mentioned above, our matching algorithm was a variant of the DA algorithm of Gale and Shapley (1962) that uses the choice function defined in Algorithm 1, and is described in Algorithm 2. This was motivated by the common view that stability of the mechanism is essential for its success and survival (see, e.g., Roth, 2002). Furthermore, when preferences are "well-behaved" (for example, responsive or substitutable), the DA is strategy-proof for the proposing side (Dubins and Freedman, 1981; Roth, 1982).[7] The preferences of the candidates (who are the proposing side in our implementation) fit exactly the scenario of Gale and Shapley (1962): a simple preference order on (a subset of) the PMAs. The preferences of the PMAs, however, are significantly more complex: beyond a preference order on the candidates, they also have the various category maximum-capacities and minimum-targets. Notably, these preferences do not necessarily satisfy substitutablity (Hatfield and Milgrom, 2005), unilateral substitutability (Hatfield and Kojima, 2010b), or even substitutable completability (Hatfield and Kominers, 2015), nor can they be described as slot-specific priorities (Kominers and Sönmez, 2016).

Can the desirable theoretic guarantees of the DA algorithm be extended to this scenario as well? Extending a known class of results (Biró et al., 2010; Huang, 2010; Fleiner and Kamiyama, 2012),

[7]Of course, that does not mean that people report their preferences truthfully, see Hassidim et al. (2017a) and references therein.

input : Set $\mathcal{M}$ of PMAs with preferences, set $\mathcal{C}$ of candidates with preferences
output: A feasible matching of (a subset of) $\mathcal{C}$ to (slots of) $\mathcal{M}$
repeat
    `// A single deferred-acceptance round`
    foreach *candidate* $c \in \mathcal{C}$ **do**
        c applies to the PMA that c ranks highest of those who have not (yet) rejected c;
    end
    foreach *PMA* m **do**
        $C \leftarrow$ the set of candidates that applied to m in this round;
        $\hat{C} \leftarrow$ the set of candidates chosen from C via m's choice function (Algorithm 1);
        m rejects all candidates in $C \setminus \hat{C}$;
    end
until *no candidate was rejected in this round*;
`// Match according to the last round above`
foreach *candidate* $c \in \mathcal{C}$ **do**
    Match c with the PMA to which it applied in the last round above;
end

Algorithm 2: Deferred-Acceptance using the choice function defined in Algorithm 1.

it turns out that this depends on the structure of overlap of the different categories. If the class of categories is *laminar* (i.e., every two categories are either disjoint or one contains the other), then, extending known results in somewhat weaker settings,[8] we show that "everything works":

Theorem 1. *If the set of categories of each of the PMAs (separately) is laminar, then Algorithm 2 produces a stable matching and is incentive compatible for the candidates.*

The key idea behind the proof of Theorem 1 is to show that the choice function defined by Algorithm 1, when categories are laminar, is *substitutable* and satisfies the *law of aggregate demand.* With these two properties in hand, Theorem 1 follows from the analysis of Hatfield and Kojima (2010a).

Unfortunately, in our setting the categories are in no sense laminar as in Theorem 1. For example, a geographic category and a religion category typically neither are disjoint nor have one containing the other. As noted above, it is commonly understood that lack of laminarity destroys desired theoretical properties of Gale-Shapley-like markets, and indeed, again extending known results in a somewhat stronger setting, we show this in a very strong sense:

Theorem 2. *If each PMA can have overlapping categories with some maximum quota (even without any minimum targets) then:*

1. *It is possible that no stable matching (w.r.t. the choice function defined in Algorithm 1) exists.*
2. *It is NP-complete to decide whether a stable matching exists or not.*
3. *Algorithm 2 is not incentive compatible for the candidates.*

[8]The most closely related previous results have considered minimum-quotas rather than minimum-targets, and have shown that laminarity implies a polynomial-time algorithm for *deciding whether* a stable matching exists. We consider minimum-targets and show for the choice function defined in Algorithm 1 that laminarity *guarantees the existence* of a stable matching and that a stable matching can be found in an *incentive-compatible* manner.

Nevertheless, based on our understanding of the problem domain as well as computer simulations, we believed that Algorithm 2 would give rather good results in practice, and this is what we set on to implement. The big question was to what extent can we evaluate the quality of this algorithm and support this belief. We now turn to discuss this, while paying special attention exactly to the properties that Theorem 2 showed we cannot fully achieve: stability and incentive compatibility.

2.2 Evaluating our Algorithm, and a Final Tweak

The level of incentive compatibility is difficult to evaluate algorithmically, and an ex-post evaluation is also less satisfying. For this we were able to formally prove that the two main types of manipulation that our candidates were considering are not profitable. Informal conversations with candidates as well as with parents of candidates while we designed the system suggested that they only considered or worried about the following two types of manipulation, each of which was brought up by quite a few candidates/parents:

- *Truncation*: many candidates were worried that listing multiple PMAs in their preference list may cause the system to match them to a less preferred option rather than a more preferred one that they would have gotten had they not allowed the system to use the less preferred option.
- *Sure thing*: many candidates knew that they have a "guaranteed slot" in a certain PMA, i.e., they were assured that the PMA did not rank above them more candidates than it has slots (in any category). Such candidates were often worried that they could lose the guaranteed slot unless they ranked this certain PMA at the top even when they really preferred another PMA.

We prove formally that neither of these two manipulations can ever be profitable.

Theorem 3. *Algorithm 2 has the following incentive properties for every candidate and every profile of preferences:*

1. *If when specifying a certain preference list, the candidate is allocated a certain outcome, then he or she will be allocated the same outcome for any extension of the original preference list.*
2. *If the candidate is guaranteed to be assigned to a certain PMA if the candidate ranks it first, where "guaranteed" means that this would happen for any profile of preference lists of the other candidates,*[9] *then if the candidate specifies any preference list that contains this PMA, the candidate is still guaranteed to either be allocated that PMA or one that the candidate placed higher on his or her preference list.*

To measure the extent of the lack of stability of an assignment (once such is reached), we turned to the standard measure of *blocking pairs*: pairs of candidate c and PMA m such that a) c prefers m over the PMA matched to c by the algorithm, and b) if m were to choose from the set of all candidates matched to her and in addition c, then the choice of m would include c. (We also used other related measures such as the number of candidates involved in blocking pairs). In simulations with various natural distributions, we have found that typically there are very few blocking pairs, which furthermore involve only a small number of candidates. (This was borne out on the real data — see below.) This may be contrasted with the outcome of the Boston mechanism

[9] And in particular if all other candidates ranked this PMA highest.

(also known as the Immediate Acceptance algorithm, see Abdulkadiroğlu et al., 2005. that does not attempt to achieve stability, and so typically results in at least an order of magnitude more of blocking pairs — giving a sneak peak into the next section, we note that on our real world data the Boston algorithm would have resulted in 147 blocking pairs, while our algorithm resulted in 10. One specific type of blocking pair that our simulations showed that Algorithm 2 produces deserves special discussion: a blocking pair that can be resolved without harming any candidate. Such a blocking pair involves a PMA m that has not fulfilled its quote and a candidate c (who may be unmatched or matched to a PMA that he or she ranked below m) such that c can be matched to m without rejecting any candidate currently matched to m (equivalently, the choice function of m, when applied to the set that includes all candidates assigned to m as well as c, chooses this entire set). Such a blocking pair can be easily resolved: simply transfer/assign c to m. Since such a resolution of this blocking pair constitutes a Pareto improvement for all candidates, then iteratively resolving any such pairs until no more such pairs exist cannot cause an infinite loop. Observing this, we heuristically added a final stage to our algorithm: a Pareto-improvement stage described in Algorithm 3. While we were able to contrive an example where Theorem 3(1) breaks

```
input  : Set 𝓜 of PMAs with preferences, set 𝓒 of candidates with preferences
output: A feasible matching of (a subset of) 𝓒 to (slots of) 𝓜
// Stage 1: Deferred acceptance
Compute a matching via Algorithm 2;
// Stage 2: Pareto improvement
while ∃ blocking pair (m, c) s.t. c can be matched to m without rejecting any candidate do
    m ← some PMA that is part of a blocking pair as defined above;
    c ← the candidate ranked highest by m s.t. (m, c) is a blocking pair as defined above;
    Match c with m (breaking any previous match of c with another PMA, if existed);
end
```

Algorithm 3: The algorithm that we ran, consisting of deferred acceptance (Algorithm 2) and an added Pareto-improvement stage.

for Algorithm 3 due to the added Pareto-improvement stage (Theorem 3(2) continues to hold even for Algorithm 3, though), we have decided to nonetheless add this stage as it seems that such examples are delicate enough so that no candidate possesses enough information as to be able to plan a successful truncation manipulation.

2.3 Epilogue

This market was run (using Algorithm 3) for the first time in January 2018. A total number of 35 different programs (administered by 24 independent PMAs) offered a cumulative 1,760 slots over which 2,580 candidates competed. The incentive properties of Theorem 3 were explained to the candidates before ranking, using a video that we prepared,[10] and in a survey conducted following the match, 92% of candidates reported that they indeed ranked truthfully. Our system matched 1,607 candidates (to 91% of available slots), 85% of whom were matched to their top choice. (Furthermore, preliminary results suggest that many of the remaining slots were quickly filled in a distributed manner following the main run of the algorithm.) The Pareto-improvement stage resolved one blocking pair: a candidate who would have been unassigned were it not for

[10] See timestamps 1:17–3:04 in `https://www.youtube.com/watch?v=xt4B2Xu3FvE`.

this stage was assigned to a PMA that previously rejected him or her. Following this stage, out of the 2580×35 possible candidate-PMA pairs, only 10 were blocking. Out of these 10, one was only theoretically problematic as it was only blocking according to our randomized tie-breaking rather than with respect to the real reported (weak) preferences. The 9 remaining blocking pairs all involved the same PMA and the same "eviction candidate" (the candidate that is to be evicted from this PMA for this blocking pair to be resolved), and all of these pairs were blocking due to one category of that PMA being one candidate short of meeting its minimum target. The candidates in all 9 of these blocking pairs were ranked by that PMA at least 2 tiers below the "eviction candidate" who was ranked at the highest tier, and so since the minimum-target of the relevant category was somewhat of a soft target ("at least 10 slots for a certain type of geographic category," of which 9 were allocated), it is very likely that the PMA would have decided to not resolve any of these blocking pairs even if it could have. The Joint Council of PMAs decided to adopt our system for the foreseeable future.

Acknowledgments

Yannai Gonczarowski is supported by the Adams Fellowship Program of the Israel Academy of Sciences and Humanities. The work of Noam Nisan is supported by ISF grant 1435/14 administered by the Israeli Academy of Sciences and by Israel-USA Bi-national Science Foundation (BSF) grant number 2014389. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 740282). The work of Assaf Romm is supported by a Falk Institute grant, by ISF grant 1780/16, and by a grant from the United States - Israel Binational Science Foundation (BSF).

We would thank the Joint Council of Pre-Military Academies, and personally Dani Zamir, Yosi Baruch, Tamar Zeira, and Yaara Schur, as well as all of the individual PMAs, for their trust in us and for their cooperation throughout the redesign process.

References

A. Abdulkadiroğlu, P. A. Pathak, A. E. Roth, and T. Sönmez. The Boston public school match. *American Economic Review*, 95(2):368–371, 2005.

A. Abdulkadiroğlu, P. A. Pathak, and A. E. Roth. Strategy-proofness versus efficiency in matching with indifferences: Redesigning the nyc high school match. *American Economic Review*, 99(5): 1954–78, 2009.

P. Biró, T. Fleiner, R. W. Irving, and D. F. Manlove. The college admissions problem with lower and common quotas. *Theoretical Computer Science*, 411(34–36):3136–3153, 2010.

L. E. Dubins and D. A. Freedman. Machiavelli and the Gale-Shapley algorithm. *The American Mathematical Monthly*, 88(7):485–494, 1981.

T. Fleiner and N. Kamiyama. A matroid approach to stable matchings with lower quotas. In *Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 135–142, 2012.

D. Fragiadakis, A. Iwasaki, P. Troyan, S. Ueda, and M. Yokoo. Strategyproof matching with minimum quotas. *ACM Transactions on Economics and Computation*, 4(1):6, 2016.

D. Gale and L. S. Shapley. College admissions and the stability of marriage. *American Mathematical Monthly*, 69(1):9–15, 1962.

I. E. Hafalir, M. B. Yenmez, and M. A. Yildirim. Effective affirmative action in school choice. *Theoretical Economics*, 8(2):325–363, 2013.

A. Hassidim, D. Marciano, A. Romm, and R. I. Shorrer. The mechanism is truthful, why aren't you? *American Economic Review*, 107(5):220–24, 2017a.

A. Hassidim, A. Romm, and R. I. Shorrer. Redesigning the Israeli psychology master's match. *American Economic Review*, 107(5):205–09, 2017b.

J. W. Hatfield and F. Kojima. Substitutes and stability for matching with contracts. *Journal of Economic Theory*, 145(5):1704–1723, 2010a.

J. W. Hatfield and F. Kojima. Substitutes and stability for matching with contracts. *Journal of Economic Theory*, 145(5):1704–1723, 2010b.

J. W. Hatfield and S. D. Kominers. Hidden substitutes. In *Proceedings of the 16th ACM Conference on Economics and Computation (EC)*, pages 37–37, 2015.

J. W. Hatfield and P. R. Milgrom. Matching with contracts. *American Economic Review*, 95(4): 913–935, 2005.

C.-C. Huang. Classified stable matching. In *Proceedings of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1235–1253, 2010.

F. Kojima. School choice: Impossibilities for affirmative action. *Games and Economic Behavior*, 75(2):685–693, 2012.

S. D. Kominers and T. Sönmez. Matching with slot-specific priorities: Theory. *Theoretical Economics*, 11(2):683–710, 2016.

T. Nguyen and R. Vohra. Stable matching with proportionality constraints. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*, pages 675–676, 2017.

N. Nisan. Bidding languages. In P. Cramton, Y. Shoham, and R. Steinberg, editors, *Combinatorial Auctions*, chapter 9, pages 400–420. MIT Press, Boston, 2006.

A. E. Roth. The economics of matching: Stability and incentives. *Mathematics of Operations Research*, 7(4):617–628, 1982.

A. E. Roth. The economist as engineer: Game theory, experimentation, and computation as tools for design economics. *Econometrica*, 70(4):1341–1378, 2002.