

Matching for the Israeli “Mechinot” Gap-Year Programs: Handling Rich Diversity Requirements

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

January 15, 2019

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, which necessitated the development of an appropriate preference-specification language along with corresponding choice-function semantics, which we also theoretically analyze to a certain extent. 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, and has been adopted by the Joint Council of the “Mechinot” gap-year programs for the foreseeable future.

1 Background

Israeli youth typically graduate from high school at the approximate age of 18 and are then required to enlist in the military 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 their military service, mostly focusing on some combination of educational and volunteering activities. It turns out that a significant number of youths, especially from higher socio-economic statuses, 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 military service that lies ahead.¹ It seems that the military also sees a benefit from these gap years, and hence it allows the participants to postpone their mandatory military service until the end of the gap year.

*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. *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.

¹The 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.

This manuscript discusses one of the most significant types of institutions that offer such gap years, called in Hebrew “Mechinot Kedam Tseva’iyot” (מְכִינֹת קֶדֶם-צְבָאִיּוֹת), in short “Mechinot” (מְכִינֹת), 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 “military-service deferments” granted by the military (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 a central part of their mission. The PMAs then build their groups incrementally 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 to give “exploding offers” to the candidates, where a candidate often has to accept an offer within days or forfeit his or her place. The candidates find it very difficult to reply to these

²See <https://mechinot.org.il/en-us/>.

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 such an exploding offer from the first 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 led to significant suboptimality 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, 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 the preferences of all of the candidates and PMAs 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 diversity or affirmative-action-type constraints

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

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

in Gale-Shapley-like matching markets. A very popular line of work considers quotas on institutions or on groups of institutions. This includes minimum quotas for specific institutions (e.g., Biró et al., 2010; Fragiadakis et al., 2016) as well as quotas on groups of institutions (e.g., Biró et al., 2010; Kamada and Kojima, 2015, 2018, and the references therein). In contrast, we are concerned not with inter-institution constraints on the number of applicants allowed in a set of institutions, but rather with intra-institution constraints for the number of applicants of a specific population within a specific institution. Such constraints in the form of minimum and maximum quotas were studied by Huang (2010) and subsequently Fleiner and Kamiyama (2012). Both of these papers consider a model with hard maximum and minimum quotas for every population in every institution. Such hard minimum quotas mean that invariably, in their model, no stable matching necessarily exists and so they focus on the study of whether a stable matching exists. Our problem differs from theirs since having had to actually build a matching mechanism, we had to output a matching even when no stable matching that perfectly matches all constraints exists. Considering affirmative action scenarios with two disjoint and complementing populations — the “majority students” and the “minority students,” Kojima (2012) use maximum quotas on the “majority students” population to effectively reserve seats for the “minority students” without explicitly using minimum quotas, thus avoiding such impossibilities. A more subtle approach to the two-population majority-minority model, which avoids some inefficiencies (such as unallocated slots) that manifest in this

מספר מקומות:		40	Number of Slots	4
מכסות מינימום/מקסימום עבור אוכלוסיות מסויימות				
שם האוכלוסיה (יש לרשום את שם האוכלוסיה במקום אוכלוסיה א/ב/...)	מינימום רצוי של משובצים מאוכלוסיה זו (ריק = אין מניעה שלא יהיו מאוכלוסיה זו)	מקסימום אפשרי של משובצים מאוכלוסיה זו (ריק = אין מניעה שכולם יהיו מאוכלוסיה זו)		
25		25	Male	8
25		25	Female	9
2		2	School	10
10		10	2 Region	11
			3 Musician	12
				13 אוכלוסיה ו
				14 אוכלוסיה ז
				15 אוכלוסיה ח

Maximum Quota

Minimum Target

Population

(a) Populations input sheet

Titles generated according to populations defined in sheet above

Musician?	Region?	School?	Female?	Male?	Ranking עדיפות לשיבוץ	ID ת.ז.	Name שם	2
		3	12	*	1	000123456	Abraham	3
*		3	14	*	1	000345678	Sarah	4
		5	17	*	2	000098765	Issac	5
		5	17	*	3	000435632	Rachel	6
*		3	12	*	3	001112223	Rebecca	7

(b) Candidate input sheet

Figure 2: A spreadsheet by which a PMA describes its preferences (with titles translated to English)

maximum-quotas approach, is taken by Hafalir et al. (2013): their approach is for each school to give priority for any minority student until the school’s “minority reserve” is filled, and beyond that to not prioritize based on populations. The requirements of the PMAs, though, involve far more than only two populations, and also involved intersecting populations. It indeed seems that the richness of the balance requirements that the PMAs desire does go significantly beyond previous analysis. A survey of the above and other approaches for handling “quotas” and diversity or affirmative-action-type constraints in stable matching markets can be found in a recent paper by Nguyen and Vohra (2017).⁴

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

⁴See also the very recent paper of Ágoston et al. (2018) for an alternative approach based on Integer Programming. While the objective that their algorithm solves for is very clear and flexible, our algorithm enjoys incentive properties (see Theorem 4 below) that were crucial in our case, and our algorithm is also somehow more transparent in its ability to answer questions along the lines of “why did candidate c not end up at PMA m ?” — questions that due to being managed independently of each other, many PMAs demanded answers to immediately following the match.

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 “populations” that it cares about (e.g., “gender,” “religiousness,” “city”). For each population the PMA is allowed to define a *maximum quota* as well as *minimum target*. In addition to ranking the candidates individually, the set of populations to which each candidate belongs is specified.⁵ Now, the interpretation of such a preference, which could be considered a significant generalization of the above-mentioned “minority reserves” approach of Hafalir et al. (2013) (in the two-population majority-minority model), 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 population never exceeds the population’s maximum quota, and strictly preferring (in a way that overrides individual ranking) candidates that belong to any population 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. It is worthwhile to mention though that in terms of cognitive simplicity, virtually none of the PMAs showed any desire to dig into these issues beyond this (non-algorithmic) general verbal interpretation, and were happy with the stability objective phrased in terms of that choice function.

Of note in Algorithm 1 is our design decision to give the same “promotion” in the preferences of a PMA to candidates that belong to any (positive) number of populations whose minimum target has not been reached, regardless of the (positive) number of such populations 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 “fill 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. This design decision required some care, though: most PMAs had a separate pool of slots for men and for women; some of these PMAs initially expressed their gender-population preferences as, e.g., “Male” population with minimum target = maximum quota = 25, and “Female” population with minimum target = maximum quota = 25. The problem with expressing their gender-population preferences this way is that its effect is to completely ignore the minimum targets of *all* populations, as any candidate who does not violate any maximum quota (and specifically, who does not violate their gender-population maximum quota) would have been given promotion since its gender minimum target would not have yet been met. For this reason, we advised these PMAs to drop the minimum targets that they set for each of the gender populations, leaving only the maximum quotas in place and resulting in proper promotion of candidates who help with “real” minimum targets.⁶ (For an alternative choice function that does more significantly promote candidates that belong to a higher number of populations whose minimum target has not been reached, see Algorithm 2.)

⁵As can be seen in Figure 2b, populations 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 population with, say, 10 possible values, is no more than a shorthand for 10 disjoint binary populations. The ability to have multiple-valued populations 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.

⁶In proof: for the matching of the following year, which is planned to take place in January 2019, we tweaked Algorithm 1 to distinguish gender populations from all other populations, so that specifically a candidate that only helps with a minimum target for a gender population is given less promotion than a candidate who helps with a minimum target for any other population, but more promotion than a candidate who helps with no minimum target.

```

input : Set  $C$  of candidates to choose from
output: Set  $\hat{C}$  of candidates chosen from  $C$ 
// Initialization
 $\hat{C} \leftarrow \emptyset$ ; // Chosen candidates
// First pass: give higher priority to candidates who help reach minimum
// targets
foreach candidate  $c \in C$ , from highest-ranked to lowest-ranked do
|   if  $c$  is in some population whose minimum quota is not met by  $\hat{C}$  then
|   |   if adding  $c$  to  $\hat{C}$  does not violate any maximum quota then
|   |   |    $\hat{C} \leftarrow \hat{C} \cup \{c\}$ ;
|   |   end
|   end
end
// Second pass: all other candidates
foreach candidate  $c \in C \setminus \hat{C}$ , from highest-ranked to lowest-ranked do
|   if adding  $c$  to  $\hat{C}$  does not violate any maximum quota then
|   |    $\hat{C} \leftarrow \hat{C} \cup \{c\}$ ;
|   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 3. 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), DA is strategyproof 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 maximum quotas and/or minimum targets on various populations. Notably, these preferences do not necessarily satisfy substitutability (Hatfield and Milgrom, 2005)⁸, 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?

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

⁸Nor of course do they satisfy unilateral substitutability (Hatfield and Kojima, 2010) or substitutable completability (Hatfield and Kominers, 2015), as both of these coincide with substitutability for matching markets without contracts, such as the PMAs market.

```

input : Set  $C$  of candidates to choose from
output: Set  $\hat{C}$  of candidates chosen from  $C$ 
// Initialization
 $\hat{C} \leftarrow \emptyset$ ; // Chosen candidates
 $\hat{R} \leftarrow \emptyset$ ; // Rejected candidates
while  $\hat{C} \cup \hat{R} \neq C$  do
    foreach candidate  $c \in C \setminus \{\hat{C} \cup \hat{R}\}$  do
        |  $n_c \leftarrow$  the number of populations of  $c$  whose minimum target is not met by  $\hat{C}$ ;
    end
     $c \leftarrow$  the highest-ranked candidate in the set  $\arg \max_{c \in C \setminus \{\hat{C} \cup \hat{R}\}} \{q_c\}$ ;
    if adding  $c$  to  $\hat{C}$  does not violate any maximum quota then
        |  $\hat{C} \leftarrow \hat{C} \cup \{c\}$ ;
    else
        |  $\hat{R} \leftarrow \hat{R} \cup \{c\}$ ;
    end
end

```

Algorithm 2: An alternative choice function for a PMA, which gives higher priority to candidates that belong to a higher number of populations whose minimum target has not been reached. (In the interest of readability, we avoided any computational speedups such as updating the values of n_c in every iteration of the while loop instead of recalculating them, keeping the candidates sorted and updating their order whenever any n_c value changes, etc..)

As noted above, Huang (2010) studies similar constraints, however in that paper the minimum constraints are hard and binding. In the absence of (positive) minimum targets/quotas, however, our model/choice function coincides with that of that paper. That paper focuses on whether or not a stable matching exists, and as that paper observes, this turns out to depend on the structure of overlap of the different populations. If the class of populations for each institution is *laminar* (i.e., every two populations are either disjoint or one contains the other), then that paper shows that in the absence of minimum quotas (and so, equivalently, in the absence of minimum targets in our model), a stable matching always exists. Furthermore, that paper shows that in the presence of minimum quotas, while laminarity obviously cannot guarantee the existence of a stable matching, laminarity does give rise to a polynomial-time algorithm for *deciding whether* a stable matching exists. In the absence of laminarity, as they show, such a decision turns out to be NP-hard. As noted above, we consider “soft” minimum targets rather than their “hard” minimum quotas. We show for the choice function defined in Algorithm 1 that laminarity of the populations, in addition to a certain condition on the structure of those populations that also have minimum targets (and not merely maximum quotas), guarantees that “everything works”: a stable matching exists, and can be efficiently found in a *strategyproof* manner.

Theorem 1. *If the set of populations of each of the PMAs (separately) is laminar, and furthermore for each PMA the populations that have (positive) minimum targets are pairwise disjoint, then Algorithm 3 (with Algorithm 1 as the choice function of the PMAs) produces a stable matching (with respect to Algorithm 1) and is strategyproof for the candidates.*

Theorem 1 can be seen as significantly generalizing aspects of the “minority reserves” approach of Hafalir et al., 2013, where there are only two disjoint populations, only one of which has a

```

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

minimum target (and so in this case the population structure is laminar and in a trivial sense the populations that have positive minimum targets are pairwise disjoint). Theorem 1 and the other theorems stated in this section are proved in Section 3. Theorem 1 is in fact a corollary of a more general theorem that we prove, which shows that for the choice function defined in Algorithm 2 (rather than that defined in Algorithm 1), a weaker condition on the structure of the populations that have minimum targets suffices.

Theorem 2. *If the set of populations of each of the PMAs (separately) is laminar, and furthermore for each PMA the set of populations that have (positive) minimum targets is a union of pairwise-disjoint chains,⁹ then Algorithm 3 with Algorithm 2 in lieu of Algorithm 1 as the choice function of the PMAs produces a stable matching with respect to Algorithm 2 and is strategyproof for the candidates.*

The key idea behind the proof of Theorem 2 (and hence Theorem 1) is to show that the choice function defined by Algorithm 1, under the assumed population structure, is *substitutable* and satisfies the *law of aggregate demand*. With these two properties in hand, Theorem 2 (as well as Theorem 1) follows from the analysis of Hatfield and Milgrom (2005). In fact, in the absence of minimum quotas (in which case Algorithm 1 and Algorithm 2 coincide with each other and with the choice function of Huang, 2010) the beginning of our proof of Theorem 2 constitutes not only a full concise alternative proof of the theorem of Huang (2010) that in the absence of any minimum quotas a stable matching exists, but also a proof that the choice function is substitutable and satisfies the law of aggregate demand, which imply many additional desirable properties (Roth, 1984; Hatfield and Milgrom, 2005) including that DA is strategyproof.¹⁰ To see that Theorem 1 indeed follows from Theorem 2, observe that a pairwise-disjoint set is a special case of a union of pairwise-disjoint

⁹Equivalently, if a population p has a (positive) minimum quota and contains two populations $q \subsetneq p$ and $r \subsetneq p$ that are incomparable (i.e., $q \setminus r \neq \emptyset$ and $r \setminus q \neq \emptyset$), then it is not the case that both q and r have (positive) minimum quotas.

¹⁰One of these additional properties is also the existence of a lattice structure for the set of stable matchings. See

chains, and that when all minimum-target populations are pairwise disjoint, Algorithms 1 and 2 coincide. As neither condition on the structure of the minimum-target populations is actually fully satisfied by the PMAs’s populations, as noted above we opted to implement Algorithm 1 as the choice function despite its slightly worse theoretical properties demonstrated above. The very few individual PMAs that did inquire more about the details of the choice function were very content with the selection of choice function, and in particular from it potentially allowing more candidates to benefit from the priority slots that they allocated.¹¹

Unfortunately, in our setting neither are the populations in any sense laminar as in Theorem 1 or Theorem 2, not do the minimum-target populations satisfy the respective structural properties. For example, a “geographic population” and a “religion population” typically neither are disjoint nor have one containing the other. As noted above, Huang (2010) has shown that even without any minimum targets/quotas, for nonlaminar populations it is possible that no stable matching exists. In this spirit, we show that dropping the condition on the structure of populations with minimum targets similarly destroys the desired theoretical properties:

Theorem 3. *If each PMA has laminar populations, however its minimum-target populations can intersect, then it is possible that no stable matching (with respect to the choice function defined in Algorithm 1) exists. Furthermore, this holds even if the set of minimum-target populations of each PMA is guaranteed to be a union of pairwise-disjoint chains.*

Nevertheless, based on our understanding of the problem domain, as well as on computer simulations, we believed that Algorithm 3 (with Algorithm 1 as the choice function of the PMAs) 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 3 showed we cannot fully achieve: stability and strategyproofness.

2.2 Evaluating our Algorithm, and a Final Tweak

The level of strategyproofness 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 and with parents of candidates while we were designing 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

Fleiner and Kamiyama (2012) for a proof, under strict minimum quotas in a generalization of the model of Huang (2010), of this property in particular.

¹¹This argument, which as noted above led us to choose Algorithm 1 as the choice function of the PMAs is indeed somewhat less convincing in a context (unlike ours) in which the set of minimum-target populations is in fact a union of pairwise-disjoint chains. Indeed, such a context it is not possible to have two intersecting minimum-target populations p and q (say, each with a minimum target of one) and two candidates c and d such that c belongs to p but not to q , and d belongs to q but not to p . Indeed, given such a population structure, if c belongs to p and d belongs to q , invariably one of c and d belongs to both p and q .

(for any population). 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.¹²

Theorem 4. *Regardless of the population structure, Algorithm 3 has the following incentive properties for every candidate and every profile of preferences, regardless of whether we use Algorithm 1 or Algorithm 2 as the choice function of the PMAs:*

- a. *If when specifying a certain preference list, the candidate is assigned to a certain PMA, then the candidate will be assigned to the same PMA even if she extends her original preference list by appending more PMAs to the end of the list.*
- b. *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,¹³ then if the candidate specifies any preference list that contains this PMA, the candidate is still guaranteed to either be assigned to that PMA or to one that the candidate placed higher on 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 it 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 (also known as the Immediate Acceptance algorithm, see Abdulkadiroğlu et al., 2005), which does not attempt to achieve stability, and so typically results in at least an order of magnitude more blocking pairs — giving a sneak peek 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 3 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 quota 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 4. While we were able to contrive an example where Theorem 4(a) breaks for Algorithm 4 due to the added Pareto-improvement stage (Theorem 4(b) continues to hold even for Algorithm 4, 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.

¹²For a formal and systematic treatment of a conceptually similar approach of ruling out “obvious manipulations,” see the recent paper by Troyan and Morrill (2018).

¹³Specifically, if the choice function of that PMA always chooses this candidate from any set of candidates that contains it.

```

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)  $\mathcal{M}$ 
// Stage 1: Deferred acceptance
Compute a matching via Algorithm 3;
// Stage 2: Pareto improvement
while  $\exists$  blocking pair  $(m, c)$  s.t.  $c$  can be matched to  $m$  without rejecting any candidate do
     $m \leftarrow$  some PMA that is part of a blocking pair as defined above;
     $c \leftarrow$  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 4: The algorithm that we ran, consisting of deferred acceptance (Algorithm 3) and an added Pareto-improvement stage.

2.3 Epilogue

This market was run (using Algorithm 4) 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 4 were explained to the candidates before ranking, using a video that we prepared,¹⁴ 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 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 population 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 population was somewhat of a soft target (“at least 10 slots for a certain type of geographic population,” 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.

3 Proofs

Proof of Theorem 2. We will show that for the assumed population structure, the choice function from Algorithm 2 is substitutable and satisfies the law of aggregate demand. By the analysis of Hatfield and Milgrom (2005), this will imply the theorem. More specifically, we will show that for every candidate set C and for every $c \in C$, if we denote by \hat{C} the set chosen from C by a given

¹⁴See timestamps 1:17–3:04 in <https://www.youtube.com/watch?v=xt4B2Xu3FvE>.

```

input : Set  $C$  of candidates to choose from
output: Set  $\hat{C}$  of candidates chosen from  $C$ 
// Initialization
 $\hat{C} \leftarrow \emptyset$ ; // Chosen candidates
 $n \leftarrow$  the number of populations with a (positive) minimum target;
for  $i \leftarrow n$  to 0 by descending order do
    foreach candidate  $c \in C$ , from highest-ranked to lowest-ranked do
        if  $c$  is in  $n$  populations whose minimum quota is not met by  $\hat{C}$  then
            if adding  $c$  to  $\hat{C}$  does not violate any maximum quota then
                 $\hat{C} \leftarrow \hat{C} \cup \{c\}$ ;
            end
        end
    end
end

```

Algorithm 5: Equivalent implementation of Algorithm 2 (along the lines of the implementation of Algorithm 1) that is traced, in parallel for two input candidate sets, in the proof of Theorem 2.

PMA and by \hat{C}' the set chosen from $C \setminus \{c\}$ by the same PMA (with the same preferences), then either $\hat{C}' = \hat{C} \setminus \{c\}$ or $\hat{C}' = (\hat{C} \setminus \{c\}) \cup \{d\}$ for some $d \in C$. This will in particular imply that $\hat{C} \setminus \{c\} \subseteq \hat{C}'$ (substitutability) and that $|\hat{C}'| \leq |\hat{C}|$ (law of aggregate demand).

In this proof we will trace the invocation of Algorithm 2, in parallel both for the input C , which we will call the “original run” of the algorithm, and for the input $C \setminus \{c\}$, which we will call the “new run” of the algorithm. To be more specific, we consider the equivalent implementation of Algorithm 2 that is given in Algorithm 5, and trace the corresponding iterations of the double-loop in both runs. We will refer to a single iteration of the inner loop as a *step* of the algorithm. These two runs have the same candidates chosen at the same steps up until candidate c is considered and chosen in the conditional in the original run. At that point in the original run candidate c is chosen, while in the new run it is not chosen (as it is not considered, since it is not in the set of candidates to choose from). From that point onward, as long as the same candidates are chosen in both runs and in the same steps, for any population the number of chosen candidates from this population in the new run is less than or equal to the same number in the old run. Therefore, either both runs accept the exact same candidates, in which case $\hat{C}' = \hat{C} \setminus \{c\}$, or the chronologically next difference in chosen candidates between both runs is that some candidate d is chosen in the new run but not in the corresponding step in the original run. We will continue the proof by induction over the number of candidates chosen (say, in the new run) until the step in which d is chosen, where the induction is in decreasing order (so formally, the induction is over $|C|$ minus this number).

Since in the original run d is not chosen at the same step in which she is chosen in the new run, this is either because she violates the maximum quota of a population to which c belongs, or she does not receive as high a promotion in the new run because adding c satisfies a minimum target for some population to which d belongs. Let us first consider the former case. In this case, adding d in the new run causes the number of chosen candidates from that maximum-quota population p to equal its maximum quota (as adding c does in the original run). So, after the step in which d is added in the new run, in neither run will candidates from population p be chosen. We claim therefore that the same candidates are chosen at the same steps in both runs from that

point onward. This can be seen by induction: for any candidate not from p (as no more candidates from p are chosen) considered at any time after d is chosen in the new run, we claim that for each population q to which this candidate belongs, the same number of candidates from q are chosen in both runs by that time. Indeed, since this candidate is not from p , by laminarity either population q is disjoint from p or it contains p , so, since by the induction hypothesis the only difference between the sets of chosen candidates in both runs by that time is that c is chosen in the original run while d is chosen in the new run, then the same number of candidates from q have been chosen in either run, so the candidate at that time is either chosen in both runs or chosen in neither run. Therefore, in this case $\hat{C}' = (\hat{C} \setminus \{c\}) \cup \{d\}$ and we are done.¹⁵ In the rest of this proof we will therefore consider the case in which d is not chosen in the original run at the same step in which she is chosen in the new run because adding c satisfies a minimum target for some population to which d belongs. In particular, we note that in this case c and d share a minimum-target population.

If $\hat{C}' = (\hat{C} \setminus \{c\}) \cup \{d\}$, then we are done. So, assume that this is not the case, and we will look at the next chronological difference (which by this assumption must exist) in acceptances after the step in which in the new run d is chosen. We reason by cases based on whether this difference is for a candidate e to be chosen in the original run but not in the corresponding step in the new run, or for a candidate e to be chosen in the new run but not in the corresponding step in the original run.

Consider first the case in which a candidate e is chosen in the original run but not (at the corresponding step at least) in the new run. Assume for contradiction that $e \neq d$. Since in the new run e is not chosen in the step in which it is chosen in the original run, this is either because it violates the maximum quota of a population to which d but not c belongs, or it did not receive as high a promotion in the new run because adding d satisfies more minimum targets for populations to which e belongs than adding c does. Either way, e shares a population with d to which c does not belong. Therefore, by laminarity of the population structure, any population that contains c and d also contains e , and any population that contains c and e also contains d . So, the minimum-target population that we have shown that c and d share also contains e . Therefore, c , d , and e share a minimum-target population, and so by the assumption on the structure of minimum-target populations, these three candidates belong to the same minimum-target population chain.

We claim that in any step from the time at which c is chosen in the original run and until the end of each run separately, d and e are in the exact same number of populations whose minimum targets are at that step not yet met. Indeed, if at any step d (resp. e) were in more such populations than e (resp. d), then since these two candidates belong to the same minimum-target population chain, this means that d (resp. e) belongs to a deeper population in that chain than e (resp. d), and that the minimum target of this “deeper” population is not yet met at that step. Since any population that contains c and d also contains e (resp. since any population that contains c and e also contains d), and since all three candidates belong to that minimum-target population chain, this means that d (resp. e) belongs to a deeper population in that chain than also c , and that the minimum target of this “deeper” population is not yet met at that step. Since the minimum target of this “deeper” population is not yet met at that step, it is also not met at any previous step. In particular, it is not met just before c is chosen in the original run (as the runs are identical before c is added), and therefore in the beginning of that step d (resp. e) is in more populations whose

¹⁵As mentioned in Section 2.1, the proof of Theorem 2 up until this point, when applied in the absence of minimum quotas (in which case Algorithm 1 and Algorithm 2 coincide with each other and with the choice function of Huang, 2010), constitutes not only a full concise alternative proof of the theorem of Huang (2010) that in the absence of any minimum quotas a stable matching exists, but also a proof that the choice function is substitutable and satisfies the law of aggregate demand, which implies many additional desirable properties (Roth, 1984; Hatfield and Milgrom, 2005) including that DA is strategyproof.

minimum target is not yet met than c is, and so d (resp. e) should have been chosen even before c , as choosing it would not have violated any maximum quotas, since d is chosen later in the new run and added to a superset of the same chosen set of candidates without violating any maximum quotas (resp. since e is later chosen in the original run). This is a contradiction, and so indeed in any step from the time at which c is chosen in the original run and until the end of each run separately, d and e are in the exact same number of populations whose minimum targets are at that step not yet met.

Since adding e does not violate any maximum quotas when it is chosen in the original run, adding it earlier in the original run, just before d is chosen in the new run, would not have violated any maximum quotas either, and so, since in the corresponding step in the new run the set of already-chosen agents is smaller, adding e in the new run at that step would not have violated any maximum quotas either. Since at that step d and e are in the exact same number of populations whose minimum targets are not yet met, and since d is chosen at that step and e is not added before even though both additions are feasible in terms of not violating any maximum quotas, we conclude that the given PMA prefers d to e .

Since e is chosen in the original run when d is not (yet) chosen, since at that step d and e are in the exact same set of populations whose minimum targets are not yet met, and since the given PMA prefers d to e , this means that choosing d at that step would have violated a maximum quota for a population to which e does not belong. Since any population to which c and d belong also contains e , this means that neither does c belong to this maximum-quota population, and so choosing d at that step (of the original run) would have violated the same maximum quota even if the set of already-chosen candidates at the time were the same set only with c removed, however this set with c removed is precisely the set of already-chosen candidates in the corresponding step in the new run, and d is chosen in that step in that run, so no maximum quotas are violated by adding d to this set — a contradiction. So, $e = d$.

Immediately following the step in which $e = d$ is chosen in the original run, the sets of chosen candidates in both runs differ only by c being only chosen in the first run. At this point we are either done if no more differences in acceptances between the two runs occur, or (by the same reasoning as in the beginning of this proof, since for any population, the number of chosen candidates from this population in the new run is less than or equal to the same number in the old run), the next chronological difference is for a candidate d' to be chosen in the new run but not in the old run. In this case by the induction hypothesis we are also done.

We now consider the case in which a candidate e is chosen in the new run but not (at that point at least) in the original run. We will complete the proof by showing that this case leads to a contradiction. Since in the original run e is not chosen in the step in which it is chosen in the new run, this is either because it violates the maximum quota of a population to which c but not d belongs, or it does not receive as high a promotion in the original run because adding c satisfies more minimum targets for populations to which e belongs than adding d does. Either way, e shares a population with c to which d does not belong. Therefore, by laminarity of the population structure, any population that contains c and d also contains e , and any population that contains d and e also contains c . So, the minimum-target population that we have shown that c and d share also contains e . Therefore, c , d , and e share a minimum-target population, and so by the assumption on the structure of minimum-target populations, these three candidates belong to the same minimum-target population chain.

We claim that in any step from the time at which c is chosen in the original run and until the end of each run separately, e is in at least as many populations whose minimum targets are at that step not yet met as d is. Indeed, if at any step d were in more such populations than e , then since they belong to the same minimum-target population chain, this means that d belongs to a deeper

population in that chain than e , and that the minimum target of this “deeper” population is not yet met at that step. Since any population that contains c and d also contains e , and since all three candidates belong to that minimum-target population chain, this means that d belongs to a deeper population in that chain than also c , and that the minimum target of this “deeper” population is not yet met at that step. Since the minimum target of this “deeper” population is not met at that step, it is also not met at any previous step. In particular, it is not met just before c is chosen in the original run (as the runs are identical before c is added), and therefore in the beginning of that step d is in more populations whose minimum target is not yet met than c is, and so d should have been chosen even before c , as choosing it would not have violated any maximum quotas, since d is chosen later in the new run and added to a superset of the same chosen set of candidates without violating any maximum quotas. This is a contradiction, and so indeed in any step from the time at which c is chosen in the original run and until the end of each run separately, e is in at least as many populations whose minimum targets are at that step not yet met as d is.

Since e is chosen after d in the new run, adding it at any time before d is chosen would not have violated any maximum quotas either. Since at that step e is in at least as many populations whose minimum targets are not yet met as d , and since d is chosen at that step and e is not added before even though both additions are feasible in terms of not violating any maximum quotas, we conclude that the given PMA prefers d to e .

Since any population that contains c and d also contains e , and any population that contains d and e also contains c , the set of populations that contain c and d is precisely the same as the set of populations that contain d and e (as both of these are precisely the same as the set of populations that contain c , d , and e).

Consider the original run in the step in which e is chosen in the new run. Adding e in this step in the new run does not violate any maximum quotas. Therefore, adding d in the same step in the original run would not have violated any maximum quotas, because any maximum quotas violated by this would have still been violated even if the set of chosen candidates at the time were the same set with c replaced by e (since the populations that contain c and d are the same as the populations that contain d and e), however in the corresponding step in the new run, this set with the addition of d is precisely the set of chosen candidates once e is chosen, and so no maximum quotas are violated by it.

Still considering the step in the which e is chosen in the new run, we now claim that the number of populations whose minimum targets are at that step in the original run not met and to which d belongs, is the same as the number of populations whose minimum targets are at that step in the new run not met and to which e belongs. To see this, we first note that any population that counts toward the former number must have c in addition to d in it (otherwise d is deeper than c in the chain of minimum-target populations to which they both belong, and furthermore the minimum target of this population is not yet met also just before c is chosen in the original run, and so d should have been chosen before that in that run), and any population that counts toward the latter must have d in addition to e in it (otherwise e is deeper than d in the chain of minimum-target populations to which they both belong, and furthermore the minimum target of this population is not yet met also just before d is chosen in the new run, and so e should have been chosen before that in that run). Therefore, since the set of populations that contain c and d is precisely the same as the set of populations that contain d and e , and both of these are precisely the same as the set of populations that contain c , d , and e , only populations that contain c , d , and e count toward any of these two numbers. Now, each of these populations can be tested to count toward the former number by counting the number of chosen candidates at that step (in the original run), except for c , that are in that population, adding one (since c is in that population), and checking if this satisfies its minimum target. Similarly, each of these populations can be tested to count toward the latter

number by counting the number of chosen candidates at that step (in the new run), except for d (this is exactly the same as the chosen candidates at that step in the original run except for c !) that are in that population, adding one (since d is in that population), and checking if this satisfies its minimum target. Therefore, each of these populations either counts toward both of these numbers or toward none of them, and so these numbers are indeed the same.

To sum up, in the step of the original run that corresponds to the step of the new run in which e is chosen, adding d would not have violated any maximum quotas, and the number of populations with unmet minimum targets at that step in the original run to which d belongs is such that d at that step (and before) passes the minimum-target conditional of the iteration of the outer loop of Algorithm 5 that contains that step. Since the given PMA prefers d to e , this means that d should have in fact already been chosen in the original run at that point — a contradiction to the assumption that the next difference between the two runs, following the addition of d in the new run, is the addition of e in the new run. \square

Proof of Theorem 3. Consider a market with three candidates c, d, e , and two PMAs m and m' . The preferences of the candidates are:

$$c : m' \succ m \quad d : m \succ m' \quad e : m \succ m'$$

The candidate ranking of PMA m is: $c \succ d \succ e$. The populations that m considers are $\{c,d,e\}$ with minimum target 1, $\{d,e\}$ with maximum quota 2, and $\{e\}$ with minimum target 1. The candidate ranking of PMA m' is: $e \succ c \succ d$. There are no populations for m' , and the overall maximum quota of m' is 1.

To see that no stable matching (with respect to the choice function defined in Algorithm 1) exists in this market, we reason by cases based on whether or not candidate c is assigned to PMA m in a proposed matching. If c is assigned to m , then since c prefers m' over m , stability dictates that m' prefer the candidate assigned to it over c . So, e is assigned to m' . But then we have that candidate e blocks with PMA m , since e prefers m over m' , and m chooses the pair $\{c, e\}$ from any candidate set that contains $\{c, e\}$ (and c is already assigned to m).

If c is not assigned to m , then since d ranks m highest and since m chooses $\{d\}$ from any candidate set that contains d but does not contain c , by stability d is alone assigned to m . Therefore, e is not assigned to m , and so by stability e is assigned to her next choice m' , since m' ranks her highest. Therefore, c is unassigned, but then we have that c blocks with m , since m chooses $\{c, d\}$ from $\{c, d\}$ (and $\{d\}$ is assigned to m). \square

Proof of Theorem 4. For Part a, note that since Algorithm 3 is a candidate-proposing deferred-acceptance algorithm, a run of this algorithm with an extended preference list for the given candidate would be identical to the original run (without the extension), as the candidate will never be rejected from the given PMA, and so the extension of the preference list will never be accessed.

For Part b, note that since Algorithm 3 is a candidate-proposing deferred-acceptance algorithm, the given candidate will be assigned to a PMA that she placed lower than the given PMA on her preference list only if the given PMA rejects her in some round (that is, does not choose her in some round in which this candidate applies to it). As it is assumed that the given PMA always chooses this candidate from any set of candidates that contains it, this cannot be, and so the candidate is either assigned to the given PMA or to one that she has placed higher on her preference list. \square

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.
- K. C. Ágoston, P. Biró, and R. Szántó. Stable project allocation under distributional constraints. *Operations Research Perspectives*, 5:59–68, 2018.
- 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, 2010.

- 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.
- Y. Kamada and F. Kojima. Efficient matching under distributional constraints: Theory and applications. *American Economic Review*, 105(1):67–99, 2015.
- Y. Kamada and F. Kojima. Stability and strategy-proofness for matching with constraints: A necessary and sufficient condition. *Theoretical Economics*, 13(2):761–694, 2018.
- 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. Stability and polarization of interests in job matching. *Econometrica*, 52(1):47–58, 1984.
- A. E. Roth. The economist as engineer: Game theory, experimentation, and computation as tools for design economics. *Econometrica*, 70(4):1341–1378, 2002.
- P. Troyan and T. Morrill. Obvious manipulations. Mimeo, 2018.