

Stability in Matching Markets with Complex Constraints*

Hai Nguyen[†] Thành Nguyen[‡] Alexander Teytelboym[§]

Abstract

We develop a model of many-to-one matching markets in which agents with multi-unit demand aim to maximize a cardinal linear objective subject to multidimensional knapsack constraints. The choice functions of agents with multi-unit demand are therefore not substitutable. As a result, pairwise stable matchings may not exist and, even when they do, may be highly inefficient. We provide an algorithm that finds a group-stable matching that approximately satisfies all the multidimensional knapsack constraints. The novel ingredient in our algorithm is a combination of matching with contracts and Scarf’s Lemma. We show that the degree of the constraint violation under our algorithm is proportional to the sparsity of the constraint matrix. The algorithm, therefore, provides practical constraint violation bounds for applications in contexts, such as refugee resettlement, daycare allocation, and college admissions with diversity requirements. Simulations using refugee resettlement data show that our approach produces outcomes that are not only more stable but also more efficient than the outcomes of the Deferred Acceptance algorithm. Moreover, simulations suggest that in practice constraint violations under our algorithm would be even smaller than the theoretical bounds.

Keywords: matching markets, market design, pairwise stability, group stability, efficiency, iterative rounding algorithms, multidimensional constraints, knapsack constraints.

JEL classification: C78, D47

*An extended abstract of this paper appeared in the Proceedings of the 2019 ACM Conference on Economics and Computation (EC’19). We would like to thank Éva Tardos, Andrew Trapp, Rakesh Vohra, participants of the 2019 Simons Institute Program on Online and Matching-Based Market Design, and three anonymous referees for their comments and suggestions.

[†]Purdue University, Department of Computer Science, Purdue University, West Lafayette, Indiana, 47906, United States. E-mail: nguye245@purdue.edu.

[‡]Krannert School of Management, Purdue University, 403 W. State Street, West Lafayette, Indiana, 47906, United States. E-mail: nguye161@purdue.edu.

[§]Department of Economics and St. Catherine’s College, University of Oxford, Oxford, OX1 3UQ, United Kingdom. This work was supported by the Economic and Social Research Council grant number ES/R007470/1. Email: alexander.teytelboym@economics.ox.ac.uk.

1 Introduction

The theory of two-sided many-to-one matching markets pioneered by [Gale and Shapley \(1962\)](#) has informed a large number of successful applications in entry-level labor markets ([Roth, 1984](#)), school choice ([Abdulkadiroğlu and Sönmez, 2003](#)), and college admissions ([Hassidim et al., 2017](#)). A key feature in the (re-)design of many successful centralized matching markets is the use of mechanisms that produce stable outcomes, i.e., matchings that are immune to blocking coalitions of agents who prefer each other to their current matches ([Roth and Xing, 1994](#); [Roth, 2002](#)).

A matching is group-stable (pairwise stable) if no group (pair) of agents who are not matched to each other would be better off by being matched together. Theory and practice offer at least three reasons why rules that produce (even pairwise) stable outcomes in matching markets perform well. First, stable allocations prevent agents' rematching to improve upon the outcome ([Roth, 1991](#)). This gives participants confidence that they do not need to waste time looking for better matches. Second, stability encodes an appealing notion of fairness in the sense of eliminating waste and justifiable envy. For example, in the context of school choice if a student misses out on a school place in a stable mechanism then no student with a lower priority could have been admitted to that school ([Abdulkadiroğlu and Sönmez, 2003](#)). Third, in large markets, stable mechanisms could also encourage participation in the marketplace and discourage early matches ([Roth and Shorrer, 2018](#)).

In classic models of many-to-one matching with maximum quotas or substitutable preferences, a pairwise stable matching exists and can be found by Gale and Shapley's Deferred Acceptance algorithm ([Gale and Shapley, 1962](#); [Kelso and Crawford, 1982](#); [Hatfield and Milgrom, 2005](#)).¹ Furthermore, in these simple settings, pairwise stable matchings coincide with group-stable matchings. Hence, pairwise stable matchings not only exist and are immune to arbitrary blocking coalitions but are also Pareto-efficient.

¹An agent's preferences satisfy substitutability if whenever an agent rejects an offer, he also rejects the offer from a larger offer set.

If agents’ preferences are not substitutable (e.g. there are complementarities or “couple”), pairwise matchings may not exist (Roth, 1984). Moreover, if preferences exhibit strong complementarities, then pairwise stable matchings may be highly inefficient. Finding group-stable matchings can mitigate efficiency losses, however, substitutability is, in fact, a necessary condition for the existence of group-stable outcomes in matching markets without transfers (Hatfield and Kojima, 2008).

In this paper, we consider a class of matching markets in which there are (i) strict, ordinal preferences for unit-demand agents (henceforth *unit agents*), (ii) linear, cardinal objective functions for agents with multi-unit demand (henceforth *multi-unit agents*), (iii) multidimensional knapsack constraints on the capacities of multi-unit agents. Our setting can, therefore, capture a very general structure of preferences and capacity constraints. In particular, unit agents can have a “size” on each dimension while multi-unit agents choose their matches by maximizing a linear objective subject to their individual multidimensional knapsack constraints. Since multi-unit agents are effectively solving their own multidimensional knapsack problems, their induced preferences can exhibit both substitutability and complementarity. In order to analyze group-stable matchings in our model, we consider fractional matchings. In a fractional matching, an agent may be matched to a fraction of another agent. One interpretation of a fractional matching is that agents are matched to each other with a given probability or for a fraction of time. We analyze the structure of fractional matchings using Scarf’s Lemma (Scarf, 1967). We provide bounds (in terms of the relative “sizes” of “largest” unit agents) on constraint violations and on the loss in the value of any multi-unit agent’s objective function which can occur when we round fractional group-stable matchings to integral group-stable matchings.

There are many applications of our model and we discuss the following three in detail (see Table 1 for a summary).

Refugee resettlement Around 50,000 refugees are resettled to the United States every

year. Nine American resettlement agencies coordinate networks of dozens of local communities (which we refer to as *localities*) with support refugees during their first few years in the United States. Localities are allowed to admit different numbers of refugees and offer different services (e.g. language support, support for large families, support for single parents etc.) that may or may not match the needs of the refugees (Delacrétaz et al., 2019; Trapp et al., 2018). In order to secure government approval to conduct resettlement, resettlement agencies must aim to maximize the number of refugees who are employed within 90 days of arrival. Some resettlement agencies have begun using integer optimization and machine learning in order to improve refugee outcomes (Bansak et al., 2018; Trapp et al., 2018), however, as far as we know, refugees’ preferences over localities are presently not collected anywhere in the world.

Daycare allocation Daycare allocation might, at first sight, appear similar to school choice: authorities aim to produce allocations that are fair and efficient with respect to preferences of parents and priorities of schools and kindergartens. Priorities are often created in order to satisfy broader societal objectives, e.g., to minimize distance walked or time traveled for children (Shi, 2015; Veski et al., 2017; Biró and Gudmundsson, 2020). As in school choice, a daycare center would typically prioritize children who live nearest to it. Therefore, the daycare center can assign a higher score to every student who lives closer (in relative or absolute terms) and then maximize the sum of the assigned scores. However, there are two crucial distinctions between school choice and daycare choice. First, families often want to send more than one child to the same daycare center because they are responsible for picking them up and dropping them off. Second, parents often want to send different numbers of children to daycare on different days of the week because they are working flexibly or part-time (Delacrétaz, 2019; Kamada and Kojima, 2019). This means that demand for a daycare center, unlike for a school, is not necessarily uniform across weekdays. Multidimensional constraints can, therefore, capture daycare capacity and parents’ demands for daycare on

different days of the week.

School and college admissions Many schools and colleges around the world take into account diversity considerations in their admissions decisions (Aygün and Turhan, 2017). Students’ diversity manifests itself on a number of dimensions e.g. income, race, or gender. A student may be in a minority category for several diversity dimensions simultaneously. A student’s type is therefore a $(0,1)$ -vector indicating membership of the relevant categories. Colleges have quotas for different categories. At the same time, instead of using simple priority lists, colleges often prefer to maximize a cardinal metric such as average test scores of admitted applicants. There are three specific examples from around the world.

- In Brazil, students applying to universities can declare their membership in any set of the disadvantaged groups (“public school”, “low income”, “black, mixed or indigenous descent”) (Aygün and Bó, 2019; Sönmez and Yenmez, 2019a).
- In Chile, the school admissions system specifies three diversity categories (“financially disadvantaged”, “special needs”, “high-achieving”) and any student can be a member of any subset (Correa et al., 2019).
- In India, students applying for publicly funded higher education institutions can declare membership of one of the constitutionally protected social categories (known as “scheduled castes”, “scheduled tribes”, and “other backward classes”) as well as other characteristics that merit affirmative action (e.g., disability) (Sönmez and Yenmez, 2019a,b). Colleges and universities, are in turn, obliged to meet complex affirmative action criteria that apply within and across these categories while admitting the most meritorious candidates.

It is worth noting that in these contexts students have “sizes” in the sense that they can take one unit of capacity in more than one diversity category (though no more than one unit on any single diversity category).

	Refugee resettlement	Daycare allocation	College admissions
Unit agents	Refugee families	Children	Students
Multi-unit agents	Localities	Daycare centers	Colleges
Objective	Max. employment	Max. distance-based score	Max. exam score
Constraints	Needs and services	Days of the week	Diversity

Table 1: Examples of applications of our model

1.1 Our contribution

Our paper makes three contributions to the theory of matching market design: modelling, structural, and algorithmic. Our modelling contribution is to develop a new theory of many-to-one fractional matching with general preferences and constraints while using a powerful solution concept. In particular, multi-unit agents in our model have preferences that arise from solving integer or linear programs. We therefore extend the theory of fractional matchings by allowing for a new preference class, multidimensional constraints, and group deviations. Matching markets with the features of our model arise naturally in a number of applications.

For our first result, we prove that fractional group-stable matchings exist using Scarf’s Lemma (Theorem 1). The main novelty here is not the existence result itself (existence can be also shown by a fixed-point argument), but developing the machinery that we apply throughout the paper. Scarf’s Lemma proves the existence of a particular *dominating vector* in a system of linear inequalities defined by a constraint matrix \mathcal{Q} . The columns of the constraint matrix describing the inequalities can be interpreted as representing possible blocking coalitions. Theorem 1 shows that such a dominating vector is close to a fractional group-stable matching. Theorem 2 then shows that with a small enough price increment, the dominating vector is a fractional group-stable matching.

We then develop an Iterative Rounding (IR) algorithm that relaxes capacity constraints in order to find an integral group-stable matching from a nearby fractional group-stable matching (Figure 1). We prove two key results. First, we show that the bound on the violation of constraints is proportional to the sparsity of the constraint matrix (in relative

terms), which is measured by its maximum column sum (Theorem 3). The maximum column sum represents the relative “size” of the “largest” unit agent. Therefore, if all unit agents are fairly small relative to the capacity of the multi-unit agents, constraint violations are likely to be minimal. Second, we provide a bound on the loss in the relative value of the objective for any multi-unit agent (Theorem 4). This bound is in terms of the relative “size” as well as the relative contribution to the objective function of the “largest” agent.

Our algorithm is a novel combination of Scarf’s Lemma and a rounding procedure. We build on the approach of Nguyen and Vohra (2018), however, we make a number of substantial extensions and modifications. The novelty of our approach is the way we apply Scarf’s Lemma to capture group stability. To use Scarf’s Lemma for stable matching problems, one needs to create a constraint matrix \mathcal{Q} , where each column corresponds to a possible blocking coalition. Hitherto, two methods to construct such a matrix have been proposed: first, with columns corresponding to all possible blocking coalitions (Scarf, 1967); second, with columns corresponding to small blocking coalitions (a hospital and a doctor, or a doctor couple and two hospitals; see Nguyen and Vohra (2018)). While Nguyen and Vohra’s method only obtains matchings that are immune to deviations by small coalitions, Scarf’s method can, in principle, capture group stability. However, the disadvantage of Scarf’s method is that the resulting constraint matrix is not sparse which prevents us from guaranteeing small error bounds in the rounding procedure.

Our main structural contribution is a new way of combining ideas from matching with contracts with Scarf’s method in order to create an appropriate constraint matrix. A contract specifies both agents in the match as well as a “price” for one of the dimensions of the knapsack constraints. Our method allows us to get the best of both methods: the sparsity of our constraint matrix does not change, but the prices specified in the contracts still allow us to capture group stability via pairwise stability.

Finally, we implement our algorithm and compare it with the traditional Deferred Acceptance algorithm using data of the refugee resettlement program run by the Hebrew Immigrant

Aid Society (HIAS) in 2017. The simulation results show that the violations of capacity constraints are much smaller in practice than the theoretical bounds. Furthermore, we show that compared with the DA algorithm our approach produces solutions that are not only more stable but also more efficient.

1.2 Related Work

Our paper contributes to several strands of the matching theory literature. Our model is closest to the recent work of [Delacrétaz et al. \(2019\)](#) and [Kamada and Kojima \(2019\)](#). There are three main differences. First, both these papers assume that multi-unit agents have strict, ordinal priority list over single-unit agents. We assume that multi-unit agents can express cardinal preferences via a linear objective function. Second, the model of [Kamada and Kojima \(2019\)](#) allows for completely general feasibility constraints whereas we, like [Delacrétaz et al. \(2019\)](#), only allow for general upper bound constraints. In particular, we do not consider minimum quotas that are, in principle, permitted by Kamada and Kojima’s model (see also [Biró et al. \(2010\)](#); [Ehlers et al. \(2014\)](#); [Fragiadakis et al. \(2016\)](#)). Third, both [Delacrétaz et al. \(2019\)](#) and [Kamada and Kojima \(2019\)](#) focus on solution concepts that are only immune to pairwise blocks (e.g. “fairness” or “weak envy-freeness”). We focus on a solution concept that rules out arbitrary coalitional deviations. We show that in our setting this group stability is necessary in order to avoid highly undesirable outcomes.

More broadly, our model complements existing models on distributional and diversity constraints in school choice (e.g., [Erdil and Kumano \(2012\)](#); [Kamada and Kojima \(2012, 2015\)](#); [Echenique and Yenmez \(2015\)](#); [Gonczarowski et al. \(2019\)](#)), matching with “sizes” (e.g., [Biró and McDermid \(2014\)](#); [Delacrétaz \(2019\)](#)), and assignment with complementarities ([Nguyen et al., 2016](#)).²

Our methods extend and adapt methods from the literature on fractional stable matchings ([Tan, 1991](#); [Roth et al., 1993](#); [Teo and Sethuraman, 1998](#); [Aharoni and Holzman, 1998](#);

²See also recent work on combinatorial clock auctions with complex knapsack constraints (e.g., [Milgrom \(2017\)](#); [Milgrom and Segal \(2020\)](#)).

Aharoni and Fleiner, 2003; Dean et al., 2006; Sethuraman et al., 2006; Király and Pap, 2008; Biró and Fleiner, 2016; Bronfman et al., 2018). Our workhorse is Scarf’s Lemma which was originally introduced in order to guarantee a non-empty core of NTU games (Scarf, 1967). A number of papers have applied Scarf’s Lemma to find (fractional) stable matchings (Király and Pap, 2008; Biró and Fleiner, 2016; Wu, 2018). Nguyen and Vohra (2018) used Scarf’s Lemma to pin down a scale-free bound on the capacity adjustment for each (and total) hospital quota that would guarantee the existence of a stable outcome in a model of doctor-hospital matching with couples. Nguyen and Vohra (2019) also consider matching problem with diversity constraints described as soft proportionality constraints (e.g., a school aims to ensure that a certain fraction of students has a disadvantaged socio-economic status). However, their simpler preference model obviated the need to worry about the sparsity of the constraint matrix. Moreover, both Nguyen and Vohra (2019) and Nguyen and Vohra (2018) were primarily concerned with pairwise stability.³

Our method complements existing approaches that deal with complex constraints in matching markets. First, one can put additional assumptions on preferences (Klaus and Klijn, 2005) or on constraints (Huang, 2010) to guarantee the existence of stable outcomes. Second, one can assume that the matching market is “large” so fractional matchings have a natural interpretation as a probability of assignment (e.g., Azevedo and Leshno (2016); Jagadeesan (2017); Che and Tercieux (2018); Noda (2018); Che et al. (2018)). Third, one can look for welfare-maximizing matchings with and without stability constraints (e.g., Irving et al. (1987); Ashlagi and Shi (2016); Trapp et al. (2018); Biró and Gudmundsson (2020)). Without further assumptions on preferences or constraints, however, the latter approach can be silent about the existence of stable outcomes. By contrast, our method requires little structure on preferences and constraints, finds stable and efficient matchings, and works in any finite market.

This paper is organized as follows. Section 2 presents our model and the definitions

³Nguyen and Vohra (2019) also consider group stability, but the preference model is simpler, and as a result, there is no need to use contracts with prices.

of stability concepts. We will use the refugee resettlement terminology of “families” and “localities” to illustrate the model. Section 2.4 gives two examples that motivate our focus on group stability, highlighting the poor performance of greedy algorithms and undesirability of pairwise stability. Section 3 presents Scarf’s Lemma and the results on the existence and properties of fractional group-stable matchings. Section 4 develops the IR algorithm that finds nearby integral group-stable matchings from a fractional group-stable matching. Section 5 presents simulation results. Section 6 is a brief conclusion. The appendix contains technical proofs.

2 Model

We adapt the model of matching with multidimensional constraints due to Delacrétaz et al. (2019) to allow for more complex objective functions for multi-unit agents and for fractional matchings.

2.1 Agents and matchings

There is a finite set $f \in F$ of unit agents that we refer to as *families* (resp. children/students) and a finite set $\ell \in L$ of multi-unit agents that we refer to as *localities* (resp. daycare centers/colleges) .

Localities provide finite capacity of different *services* (resp. days of the week/diversity categories) from a finite set S . We normalize the capacity of each service at each locality to 1. We often refer to these capacities as *service constraints*. Each family requires some fraction of the capacity of each service at each locality. Let $\alpha_{f,\ell}^s \in [0, 1]$ denote the fraction of the capacity of service s that family f requires at ℓ . Note that we allow families to require some services in some localities but not in others. Several families can be matched to a locality, but any family can be matched to at most one locality. For each family, there is a set of *acceptable localities* $\bar{L}_f \subseteq L$ and for each locality ℓ there is a set of *acceptable families* $\bar{F}_\ell \subseteq F$. Consider

the set of *mutually acceptable* family-locality pairs $A = \{(f, \ell) \mid \ell \in \bar{L}_f \text{ and } f \in \bar{F}_\ell\}$.

Definition 1. A vector $\mu \in \mathbb{R}^{|F| \times |L|}$ is a *fractional matching* if

- $\mu_{f,\ell} = 0$ for all $(f, \ell) \in (F \times L) \setminus A$,
- $0 \leq \mu_{f,\ell} \leq 1$ for all $(f, \ell) \in A$,
- $\sum_{\ell \in L} \mu_{f,\ell} \leq 1$ for all $f \in F$, and
- $\sum_{f \in F} \alpha_{f,\ell}^s \mu_{f,\ell} \leq 1$ for all $\ell \in L$ and $s \in S$.

Let us examine the definition of fractional matchings. We can think of μ as a vector of *matching probabilities*. The first two conditions say that any family-locality pair that is not mutually acceptable must have a matching probability of zero, but any mutually acceptable pair can have a positive matching probability. The third condition says that the total probability of assignment across localities for any given family cannot be greater than 1. Therefore, the remaining probability $1 - \sum_{\ell \in L} \mu_{f,\ell}$ is the family's probability of being unmatched. The final condition is a feasibility condition: it says that the probability-weighted service needs of the families matched to a locality cannot exceed the locality's capacity for any service.

For a fractional matching μ , denote by L_f^μ (resp. F_ℓ^μ) the set of localities (resp. families) for which $\mu_{f,\ell} > 0$. For any fractional matching μ and any f define a vector $\mu_f \in \mathbb{R}^{|L|}$ such that $\mu_f = (\mu_{f,\ell})_{\ell \in L}$. For any fractional matching μ and any ℓ define a vector $\mu_\ell \in \mathbb{R}^{|F|}$ such that $\mu_\ell = (\mu_{f,\ell})_{f \in F}$. A fractional matching μ is an *integral matching* if $\mu_{f,\ell} \in \{0, 1\}$.

2.2 Preferences

Families have preferences over localities. We denote by \succ_f the strict ordinal preference list of family f over \bar{L}_f (localities $L \setminus \bar{L}_f$ are unacceptable for f). We write $\ell \succeq_f \ell'$ (f *weakly prefers* ℓ to ℓ') to denote that either $\ell \succ_f \ell'$ (f *prefers* ℓ to ℓ') or $\ell = \ell'$.

Families' preference lists also allow families to compare some fractional matchings.

Definition 2. A family f is *missing out* on an acceptable locality $\ell \in \bar{L}_f$ under a fractional matching μ if

- $\sum_{\ell' \in L} \mu_{f,\ell'} = \sum_{\ell' \in \bar{L}_f} \mu_{f,\ell'} < 1$, or
- there exists a locality $\ell' \in \bar{L}_f$ such that $\ell \succ_f \ell'$ and $\mu_{f,\ell'} > 0$.

Given family f 's preference list, there are two ways in which f can miss out on an acceptable locality ℓ . First, f might want to shift its probability mass from being unmatched to ℓ . If f has some probability of being unmatched, it will be missing out on all its acceptable localities \bar{L}_f . Alternatively, even if f has a zero probability of being unmatched, f might prefer to shift its probability mass from a less preferred acceptable locality ℓ' where it has a positive probability of assignment to the more preferred acceptable locality ℓ .

Preferences of each locality arise from a maximizing a linear objective subject to a set of multidimensional knapsack constraints. In particular, each locality ℓ assigns a *value* $v_{f,\ell} \in \mathbb{R}_{\geq 0}$ to each $f \in \bar{F}_\ell$ ($v_{f,\ell} = -\infty$ for any $f \in F \setminus \bar{F}_\ell$). Given a set of available families (i.e. the locality's choice set) $F' \subseteq F$, locality ℓ solves the following integer program $\text{INT}_\ell(F')$

$$\begin{aligned}
& \text{maximize} && \sum_{f \in F'} v_{f,\ell} \cdot x_{f,\ell} \\
& \text{subject to} && \sum_{f \in F'} \alpha_{f,\ell}^s \cdot x_{f,\ell} \leq 1 \quad \text{for every } s \in S && (\text{INT}_\ell(F')) \\
& && x_{f,\ell} \in \{0, 1\} \quad \text{for every } f \in F'.
\end{aligned}$$

Denote by $V_\ell(F')$ the *value of objective* $\text{INT}_\ell(F')$. We say that a locality prefers (resp. weakly prefers) F' to F if $V_\ell(F') > V_\ell(F)$ (resp. $V_\ell(F') \geq V_\ell(F)$).

We also consider a relaxation of $\text{INT}_\ell(F')$ to the following linear program $\text{LIN}_\ell(F')$

$$\begin{aligned}
& \text{maximize} && \sum_{f \in F'} v_{f,\ell} \cdot x_{f,\ell} \\
& \text{subject to} && \sum_{f \in F'} \alpha_{f,\ell}^s \cdot x_{f,\ell} \leq 1 \quad \text{for every } s \in S && (\text{LIN}_\ell(F')) \\
& && x_{f,\ell} \in [0, 1] \quad \text{for every } f \in F'.
\end{aligned}$$

Denote by $\tilde{V}_\ell(F')$ the *value of objective* $\text{LIN}_\ell(F')$.

2.3 Stability concepts

In this paper, we use a several strong stability concepts in the context of fractional and integral matchings.

Definition 3. An integral matching μ is *group-stable* if for any locality $\ell \in L$ and for any set of families $F' \subseteq F$ that are missing out on ℓ under μ , we have that $V_\ell(F_\ell^\mu) \geq V_\ell(F_\ell^\mu \cup F')$.

An integral matching is group-stable if there is no *blocking coalition* i.e locality ℓ and group of families F' such that all families F' prefer ℓ to their current matches and ℓ prefers to rematch with all families F' (perhaps alongside some families in F_ℓ^μ). Note that since $V_\ell(\cdot)$ is monotonic, $V_\ell(F_\ell^\mu) \geq V_\ell(F_\ell^\mu \cup F')$ actually implies that $V_\ell(F_\ell^\mu) = V_\ell(F_\ell^\mu \cup F')$.

Our next definition extends the definition of group stability to fractional matchings.

Definition 4. A fractional matching μ is *group-stable* if for any locality $\ell \in L$ and for any set of families $F' \subseteq F$ that are missing out on ℓ under μ , we have that $\tilde{V}_\ell(F_\ell^\mu) \geq \tilde{V}_\ell(F_\ell^\mu \cup F')$.

A fractional matching is also group-stable if there is no *blocking coalition* i.e. a locality ℓ and a group of families F' who are missing out on ℓ would prefer to change their matching probabilities. More precisely, for every locality ℓ , μ attains the maximum value of the objective even when we expand the choice set of ℓ to include any subset of the families who missing out on ℓ . Notice, this definition actually implies $\tilde{V}_\ell(F_\ell^\mu) = \tilde{V}_\ell(F_\ell^\mu \cup F')$ as $\tilde{V}(\cdot)$ is monotonic.

We also need the following notion of approximate group stability.

Definition 5. A fractional matching μ is ε -*group-stable* (for some $\varepsilon > 0$) if for any locality $\ell \in L$ and for any set of families $F' \subseteq F$ that are missing out on ℓ under μ , we have that $V_\ell(F_\ell^\mu) \geq \tilde{V}_\ell(F_\ell^\mu \cup F') - \varepsilon$.

A fractional matching is ε -group-stable if any ℓ could improve the current value of its objective (under μ) by at most ε when its choice set is expanded to include families that are missing out on ℓ .

Finally, we define both fractional and integral pairwise stable matchings which are immune to blocking coalitions comprising at most one locality and one family.

Definition 6. A fractional (resp. integral) matching is *pairwise stable* if for any locality $\ell \in L$ and for any family $f' \in F$ that is missing out on ℓ under μ , we have that $\tilde{V}_\ell(F_\ell^\mu) \geq \tilde{V}_\ell(F_\ell^\mu \cup \{f'\})$ (resp. $V_\ell(F_\ell^\mu) \geq V_\ell(F_\ell^\mu \cup \{f'\})$).

In a pairwise stable matching, there is no locality ℓ and family f' such that f' prefers to ℓ to its current match and ℓ prefers to rematch with f' (while possibly keeping some of its current families).

2.4 Inadequacy of pairwise stability

In classical settings of many-to-one matching, group-stable and pairwise stable matchings coincide and both are Pareto-efficient. This is no longer the case for our model. In a setting of matching with multidimensional knapsack constraints, pairwise stable outcome can be extremely inefficient. We now provide two examples that illustrate this point.⁴

Example 1. Suppose there is one locality ℓ and one service s . There is a set of families $\{f_1, \dots, f_{101}\}$. Suppose that $\alpha_{f_1}^s = 1$ and $\alpha_{f_i}^s = \frac{1}{100}$ for all $i \in \{2, \dots, 101\}$. Suppose that $v_{f_1} = 1$ and $v_{f_i} = 0.99$. Then, there is pairwise stable matching in which (i) only f_1 is matched to ℓ , and (ii) the $V_\ell(\{f_1\}) = 1$. This matching is, however, not a group-stable matching. Locality ℓ would prefer to block this outcome with any group of agents (of cardinality 2 or larger) from $\{f_2, \dots, f_{101}\}$. The unique group-stable matching matches one hundred families f_i for $i \in \{2, \dots, 101\}$ to ℓ yielding $V_\ell(\{f_2, \dots, f_{101}\}) = 99$. \square

⁴Since there is only one locality in both examples, we abuse notation by suppressing the ℓ subscript thereby using α_f^s to denote $\alpha_{f,\ell}^s$ and v_f to denote $v_{f,\ell}$.

The first example emphasizes that very inefficient pairwise stable matching might potentially exist. Nevertheless, when $|S| = 1$, appropriate greedy algorithms for the knapsack problem (that adjust value of the family by its “size”) are unlikely to result in such highly inefficient outcomes. However, greedy algorithms might still be very inefficient when $|S| > 1$ as the following example shows.⁵

Example 2. Suppose there is one locality ℓ and two services s_1 and s_2 . There is a set of families $\{f_1, f_2, \dots, f_{102}\}$. Suppose that $\alpha_{f_1}^{s_1} = \alpha_{f_2}^{s_2} = 1$, $\alpha_{f_1}^{s_2} = \alpha_{f_2}^{s_1} = 1/100$, $\alpha_{f_i}^{s_j} = \frac{1}{100}$ for all $i \in \{3, \dots, 102\}$ and $j \in \{1, 2\}$. Suppose that $v_{f_1} = v_{f_2} = 3$ and $v_{f_i} = 1$ for all $i \in \{3, \dots, 102\}$. Then the greedy algorithm which ranks families according to the adjusted value $\tilde{v}_f = v_f(\frac{1}{\alpha_f^{s_1}} + \frac{1}{\alpha_f^{s_2}})$ only matches f_1 yielding $V_\ell(\{f_1\}) = 3$. This is because the adjusted values are: $\tilde{v}_{f_1} = \tilde{v}_{f_2} = 3 \times (\frac{100}{1} + \frac{100}{100}) = 303$ and $\tilde{v}_{f_i} = 1 \times (\frac{100}{1} + \frac{100}{1}) = 200$ for all $i \in \{3, \dots, 102\}$. However, the group-stable matching matches 100 families $i \in \{3, \dots, 102\}$ yielding $V_\ell(\{f_3, \dots, f_{102}\}) = 100$. \square

We now show that group stability overcomes undesirable properties of a pairwise stable matchings and yields more efficient matchings.

An integral matching is wasteful if an unmatched group of families can strictly improve the objective of some locality that is acceptable to all these families.

Definition 7. An integral matching μ is *wasteful* if there is a non-empty set of families $F' \subseteq F$ and a locality ℓ such that

- $\ell \in \bar{L}_f$ for all $f \in F'$,
- $\sum_{\ell \in L} \mu_{f,\ell} = 0$ for all $f \in F'$,
- $V_\ell(F' \cup F_\ell^\mu) > V_\ell(F_\ell^\mu)$.

Proposition 1. *Suppose that $A = F \times L$, that is, all localities and families are mutually acceptable. If μ is group-stable then it is not wasteful.*

⁵Similar examples are given in [Noda \(2018\)](#) and [Ashlagi et al. \(2020\)](#).

Proof. Suppose that μ is wasteful. Consider F' and ℓ as in the definition of wastefulness. Since all family-locality pairs are acceptable then all families in F' must be missing out on every locality. In particular, all families in F' are missing out on ℓ . But ℓ prefers $F' \cup F_\ell^\mu$ to F_ℓ^μ . Hence, μ is not group-stable. \square

We now show that group stability implies Pareto optimality from the point of view of families.

Definition 8. An integral matching μ is *Pareto-dominated for families* by another integral matching μ' if for all $f \in F$ we have that $L_f^{\mu'} \succeq_f L_f^\mu$ and for at least one family $L_f^{\mu'} \succ_f L_f^\mu$.

Proposition 2. *If a group-stable integral matching μ is Pareto-dominated for families by another integral matching μ' , then all localities weakly prefer μ to μ' .*

Proof. Suppose some locality ℓ prefers $F_\ell^{\mu'}$ to F_ℓ^μ . Therefore $F_\ell^{\mu'} \neq F_\ell^\mu$. Consider all families that switched their localities in μ to locality ℓ in μ' . All these families prefer ℓ to their localities in μ . This implies that μ was not group-stable. \square

Note that in the previous result, we do not require that μ' is group-stable. Therefore, there is no efficiency loss when considering group-stable outcomes. The preceding two propositions give us the following corollary.

Corollary 1. *Given a group-stable integral matching μ , there cannot be any other matching μ' weakly preferred by all families such that $\sum_{\ell \in L} V_\ell(F_\ell^{\mu'}) > \sum_{\ell \in L} V_\ell(F_\ell^\mu)$.*

Therefore, given any group-stable matching μ we cannot find another matching μ' (group-stable or otherwise) that is weakly preferred by all families and that yields a higher total value of localities' objectives.

3 Fractional group-stable matchings

In order to eventually obtain an integral group-stable matching, we first use Scarf's Lemma to construct a fractional group-stable matching. We then use a rounding procedure to obtain

an integral group-stable matching by relaxing some service constraints. At a high level, our approach is similar to that of [Nguyen and Vohra \(2018\)](#). However, unlike [Nguyen and Vohra \(2018\)](#), we are interested in group-stable matchings and we have multidimensional knapsack constraints so we develop a novel way of using Scarf’s Lemma.

3.1 Scarf’s Lemma

In order state Scarf’s Lemma, we need the following definition which is closely related to various stability concepts.

Definition 9. Let \mathcal{Q} be an $n \times m$ non-negative matrix and let $q \in \mathbb{R}_+^n$ be a strictly positive vector. Associated with each row $i \in \{1, \dots, n\}$ of \mathcal{Q} is a weak order \succeq_i over the set of columns j for which $\mathcal{Q}_{i,j} > 0$. A vector $x \geq 0$ satisfying $\mathcal{Q}x \leq q$ *dominates* column j of \mathcal{Q} if there exists a row i such that $\sum_{j=1}^m \mathcal{Q}_{i,j}x_j = q_i$ and $k \succeq_i j$ for all $k \in \{1, \dots, m\}$ such that $\mathcal{Q}_{i,k} > 0$ and $x_k > 0$. In this case, we also say x *dominates column j at row i* .

To interpret this definition it is helpful to reinterpret matrix \mathcal{Q} whose entries are in the interval $[0, 1]$. Let us associate each row of \mathcal{Q} with an agent and interpret each column as the characteristic vector of a coalition of agents. Thus, $\mathcal{Q}_{i,j} > 0$ means that agent i is in the j^{th} coalition. The weak order \succeq_i can be seen as agent i ’s weak preference ordering over all the columns/coalitions of \mathcal{A} that contains agent i . Each non-zero component of a feasible vector x , in other words $\mathcal{Q}x \leq q$, corresponds to a coalition. For the domination, consider a coalition, say S , not selected by x . If x dominates S , there is at least one agent $i \in S$ such that i weakly prefers all the coalitions in the support of x (i.e., that include it) over S .

We use the following version of Scarf’s Lemma due to [Király and Pap \(2008\)](#).

Lemma 1 ([Scarf, 1967](#); [Király and Pap, 2008](#)). *Let \mathcal{Q} be an $n \times m$ non-negative matrix and $q \in \mathbb{R}_+^n$ be a strictly positive vector. Then there exists an extreme point of $\{x \in \mathbb{R}_+^m : \mathcal{Q}x \leq q\}$ that dominates every column of \mathcal{Q} .*

Throughout the paper we will refer to x as a *dominating vector*.

3.2 Intuition for our approach

Nguyen and Vohra (2018) use Scarf’s Lemma by interpreting each column of the constraint matrix as a possible blocking coalition, which in their case is either a pair of a doctor and a hospital or a small group containing a couple of doctors and two hospitals. A dominating vector in this context is therefore only immune to blocking coalitions of small groups of agents.

So how can we capture *all* blocking coalitions? One possible way to do this is to expand the columns of the constraint matrix \mathcal{Q} to allow for all possible blocking coalitions (see, for example, Scarf (1967) and Wu (2018)). Unfortunately, the sparsity of the constraint matrix, measured by the maximum column sum, would significantly increase, which would not allow us to obtain a satisfactory bound on constraint relaxation required to round from a fractional group-stable matching to an integral group-stable matching.

The main technical innovation in this paper is to reinterpret the constraint matrix \mathcal{Q} in a novel way. First, we associate each column of \mathcal{Q} with a family-locality pair. We then expand the set of columns in order to capture *contract terms* between the family and the locality rather than trying to capture larger coalitions. The value of an acceptable family f at a locality ℓ , i.e. $v_{f,\ell}$, can be thought of as a “budget” for the match. The contract between f and ℓ specifies how this budget is “spent” among the service constraints at ℓ . We then create “preferences” for the families and for the service constraints over the contracts, based on how the budget in the contracts is spent on these constraints.

We show that in this case a dominating vector corresponds to a group-stable matching. In particular, the way in which the budget is spent in each budget at the dominating vector can be interpreted as “competitive prices” for each service. This interpretation allows us to use the Complementary Slackness Theorem from linear programming (see Appendix A) to show that in the extended model of matching with contracts, a pairwise stable matching maps back to a group-stable matching in the original matching problem.

3.3 Applying Scarf's Lemma to matching markets with complex constraints

We now formally show how to apply Scarf's Lemma to the problem of matching with multidimensional knapsack constraints.

For each locality ℓ and each family f , we associate each service constraint s at ℓ with a price $p_{f,\ell}^s$. This price is the amount that family f would like to spend on the service constraint s at locality ℓ . Price $p_{f,\ell}^s$ takes values from a finite set $\{0, \epsilon, 2\epsilon, \dots\}$, where $\epsilon > 0$ is a price increment which can be chosen arbitrarily.

A contract between family f and locality ℓ is a variable $x_{f,\ell,p}$ which can take values between 0 and 1 where p is a price vector of length $|S|$ specifying one price $p_{f,\ell}^s$ for each service s at ℓ . Contract $x_{f,\ell,p}$ can therefore be interpreted as a probability that f is matched to ℓ at a price vector p . Recall that we can interpret $v_{f,\ell}$ as the budget for the match between f and ℓ . To avoid "over-spending" the budget, we assume that

$$\sum_{s \in S} \alpha_{f,\ell}^s \cdot p_{f,\ell}^s \leq v_{f,\ell} \quad (1)$$

for all f and ℓ . Let $P_{f,\ell}$ be the set of all feasible price vectors p satisfying equation (1).

Let $x_{f,\ell,p} \in [0, 1]$ be a vector that solves the following system

$$\sum_{f \in F} \sum_{p \in P_{f,\ell}} \alpha_{f,\ell}^s \cdot x_{f,\ell,p} \leq 1 \quad \text{for every } \ell \in L \text{ and } s \in S, \quad (2)$$

$$\sum_{\ell \in L} \sum_{p \in P_{f,\ell}} x_{f,\ell,p} \leq 1 \quad \text{for every } f \in F. \quad (3)$$

Define

$$x_{f,\ell} := \sum_{p \in P_{f,\ell}} x_{f,\ell,p} \text{ for all } f, \ell. \quad (4)$$

Then the vector $(x_{f,\ell})$ is, in fact, a fractional matching.

Our goal is to use Scarf's Lemma for the system (2)–(3), and show that the fractional

matching $x_{f,\ell}$ is group-stable. Let \mathcal{Q} be the matrix whose entries are coefficients ($\alpha_{f,\ell}^s$ and 1) in the system (2)—(3). The rows of \mathcal{Q} correspond to each family and to each locality-service pair. Hence, the number of rows of \mathcal{Q} is $|L| \cdot |S| + |F|$. Each column of \mathcal{Q} corresponds to a contract $x_{f,\ell,p}$ between f and ℓ and a price vector p . Hence, there are $\sum_{f,\ell} |P_{f,\ell}|$ columns in \mathcal{Q} .

We then define an ordering of each row of \mathcal{Q} over the columns that are in the support of that row. As in Scarf's Lemma, column j in the support of row i if $\mathcal{Q}_{ij} > 0$. The ordering is defined as follows:

- The ordering of a row corresponding to a family f is based on the preference list \succ_f of family f over localities. The row corresponding to a family f is indifferent between any two columns $x_{f,\ell,p}$ and $x_{f,\ell,p'}$ for any $p \neq p'$.
- The ordering of a row corresponding to locality-service pair (ℓ, s) is based on the decreasing order of the price $p_{f,\ell}^s$. Therefore, the (ℓ, s) -row prefers $x_{f,\ell,p}$ to $x_{f',\ell,p'}$ if $p_{f,\ell}^s > p_{f',\ell}^{s'}$. Therefore, (ℓ, s) -row prefers the contract that “pays” a higher price for s . The (ℓ, s) -row is indifferent between the other columns.

By Scarf's Lemma, there exists an extreme point x^* of $\{x : \mathcal{Q}x \leq 1\}$ that dominates every column of \mathcal{Q} .

Theorem 1. *For any price increment $\epsilon > 0$, let $(x_{f,\ell,p}^*)$ be a dominating vector of (2)—(3). Then the vector $(x_{f,\ell}^*)$ defined as in (4) is an ϵ -group-stable matching, where $\epsilon = \epsilon\Delta$, $\Delta = (1 + \alpha)|S|$, and $\alpha = \max_{f,\ell} \sum_{s \in S} \alpha_{f,\ell}^s$.*

Proof of Theorem 1. By the dominating property for every column (f_0, ℓ_0, p_0) one of the following cases must happen:

1. For the row corresponding to family f_0 the constraint binds, that is, $\sum_{\ell \in L} x_{f_0,\ell}^* = 1$ and for any locality ℓ' such that $\ell \succ_{f_0} \ell'$ there is no contract between f_0 and ℓ' , that is, $x_{f_0,\ell'}^* = 0$.

2. There is a service s_0 such that $\alpha_{f_0, \ell_0}^{s_0} > 0$ and the corresponding constraint binds, that is, $\sum_{f \in F} \sum_{p \in P_{f, \ell}} \alpha_{f, \ell_0}^{s_0} \cdot x_{f, \ell_0, p}^* = 1$. Furthermore, row (ℓ_0, s_0) weakly prefers all contracts $x_{f, \ell_0, p}^* > 0$ to contract x_{f_0, ℓ_0, p_0} .

Fix a locality ℓ , consider $F' \subseteq F$ to be the set of families that are missing out on locality ℓ under matching $x^* := (x_{f, \ell, p}^*)$. To prove that x^* is a $\Delta\epsilon$ -group-stable matching, we need to show that $\tilde{V}(F_\ell^{x^*}) \geq \tilde{V}(F_\ell^{x^*} \cup F') - \Delta\epsilon$.

Note that the dual (D) of the linear program $\text{LIN}_\ell(F'')$ is:

$$\begin{aligned} & \text{minimize} && \sum_{s \in S} p_\ell^s + \sum_{f \in F''} q_f \\ & \text{subject to} && q_f + \sum_{s \in S} \alpha_{f, \ell}^s \cdot p_\ell^s \geq v_{f, \ell} \quad \text{for every } f \in F'', \\ & && p_\ell^s \geq 0, q_f \geq 0 \quad \text{for every } s \in S, f \in F''. \end{aligned}$$

Observe that the dual variables p_ℓ^s and q_f correspond to the primal constraints of the linear program $\text{LIN}_\ell(F'')$. Let price vector \bar{p}_ℓ be the coordinate-wise minimum over the set of all price vectors p such that $x_{f, \ell, p}^* > 0$ for some family f . And let $\bar{q}_f = 0$ if the corresponding primal constraint is not binding (i.e., if $x_{f, \ell} < 1$) and $\bar{q}_f = v_{f, \ell} - \sum_{s \in S} \alpha_{f, \ell}^s \cdot \bar{p}_\ell^s$ otherwise. We now state two key claims that we prove in Appendix B.1.

Claim 1. *If $x_{f, \ell}^* > 0$, then $\bar{q}_f + \sum_{s \in S} \alpha_{f, \ell}^s \cdot \bar{p}_\ell^s \geq v_{f, \ell} - \epsilon \sum_{s \in S} \alpha_{f, \ell}^s$.*

Claim 2. *If $\sum_{f \in F} \alpha_{f, \ell}^s \cdot x_{f, \ell}^* < 1$, then $\bar{p}_\ell^s \leq \epsilon \sum_{s \in S} \alpha_{f, \ell}^s$, and if $x_{f, \ell}^* < 1$ then $\bar{q}_f = 0$.*

Intuitively, the first claim says that if the primal variable is slack, then the corresponding dual constraint must be binding when ϵ goes to 0. The second claim says that if the primal constraint is slack, then the corresponding dual variable must be zero when ϵ goes to 0.

By Claims 1, 2, and the way \bar{q}_f was chosen, we have

$$\left(\bar{q}_f + \sum_{s \in S} \alpha_{f, \ell}^s \cdot \bar{p}_\ell^s - v_{f, \ell} \right) \cdot x_{f, \ell}^* \leq 0 \text{ for every family } f, \quad (5)$$

$$\left(1 - \sum_{f \in F} \alpha_{f,\ell}^s \cdot x_{f,\ell}^*\right) \cdot \bar{p}_\ell^s \leq \left(1 - \sum_{f \in F} \alpha_{f,\ell}^s \cdot x_{f,\ell}^*\right) \cdot \epsilon \cdot \alpha, \quad (6)$$

$$(1 - x_{f,\ell}^*) \bar{q}_f = 0. \quad (7)$$

Now adding (7), (8), and (9) over all possible f and s ,

$$\begin{aligned} \sum_s \bar{p}_\ell^s + \sum_f \bar{q}_f - \sum_f v_{f,\ell} x_{f,\ell}^* &\leq \epsilon \alpha \left(\sum_s 1 - \sum_s \sum_f \alpha_{f,\ell}^s \cdot x_{f,\ell}^* \right) \quad (\text{for every service } s) \\ &\leq \epsilon \alpha \sum_s 1 \\ &= \epsilon \alpha |S|. \end{aligned}$$

This implies that

$$\sum_f v_{f,\ell} x_{f,\ell}^* \geq \sum_s \bar{p}_\ell^s + \sum_f \bar{q}_f - \epsilon \alpha |S|.$$

Note that $(\bar{p}_\ell, \bar{q}_f)$ is not a feasible solution for the dual but $(\hat{p}_\ell, \bar{q}_f)$ is a feasible solution, where $\hat{p}_\ell = \bar{p}_\ell + \epsilon$. Hence,

$$\sum_f v_{f,\ell} x_{f,\ell}^* \geq \sum_s \hat{p}_\ell^s + \sum_f \bar{q}_f - \epsilon |S| - \epsilon \alpha |S| = \sum_s \hat{p}_\ell^s + \sum_f \bar{q}_f - \epsilon(1 + \alpha)|S|.$$

Therefore, the dominating vector x^* is $\Delta\epsilon$ -group-stable matching. \square

Theorem 1 shows that the vector $(x_{f,\ell}^*)$ obtained from the dominating vector is a fractional ϵ -group-stable matching. Theorem 1 does not, however, provide a finite algorithm to find an exact fractional group-stable matching. Our next theorem shows that we obtain an exact fractional group-stable matching if the price increment ϵ is chosen to be small enough.

Theorem 2. *There exists $\epsilon^* > 0$, such that if the price increment is $\epsilon < \epsilon^*$, then the matching $(x_{f,\ell}^*)$ obtained from the dominating vector $(x_{f,\ell,p}^*)$ in (4) is a fractional group-stable matching.*

The proof of Theorem 2 is surprisingly simple so we provide it immediately.

Proof. According to Lemma 1, a dominating vector is a vertex of the polytope defined by (2)—(3). Hence, the matching $(x_{f,\ell}^*)$ is also a vertex of the following matching polytope:

$$\begin{aligned} \sum_{f \in F} \alpha_{f,\ell}^s \cdot x_{f,\ell} &\leq 1 && \text{for all } \ell \in L, s \in S, \\ \sum_{\ell \in L} x_{f,\ell} &\leq 1 && \text{for all } f \in F. \end{aligned}$$

Let X be the set of these vertices that do not correspond to a fractional group-stable matching. There is a finite number of vertices, hence X is a finite set. We want to show that when ϵ is small enough, the matching $(x_{f,\ell}^*)$ cannot be one of the vertices in X .

If a vector $x \in X$ is not a fractional group-stable matching, then there exists $\epsilon_x > 0$ such that x is *not* a ϵ_x -group-stable fractional matching. Let

$$\epsilon_{min} := \min_{x \in X} \epsilon_x > 0.$$

According to Theorem 1 for a price increment ϵ the fractional matching $(x_{f,\ell}^*)$ is a vertex and is $\Delta\epsilon$ -group-stable. If $\Delta\epsilon < \epsilon_{min}$, then this dominating vector cannot belong to X . Hence, $(x_{f,\ell}^*)$ is a fractional group-stable matching. \square

4 Near-feasible group-stable integral matchings

Armed with a fractional group-stable matching which we obtained via Scarf's Lemma, we can now state our main result.

Theorem 3. *There exists a polynomial time algorithm that for any fractional group-stable matching finds an integral group-stable matching in which the total number of matched fam-*

ilies does not decrease if any service capacity of any locality can be increased by at most

$$\max_{f,\ell} \left(\sum_{s \in S} \alpha_{f,\ell}^s \right).$$

Proof. See Appendix B.2 □

Theorem 3 says that from any fractional group-stable matching we can obtain an integral group-stable matching that only mildly violates capacity constraints. Recall that $\alpha_{f,\ell}^s$ is the fraction of the capacity of service s that family f requires at ℓ . Therefore, the maximum column sum represents the relative “size” of the “largest” unit agent. Therefore, if all unit agents are fairly small relative to the capacity of the multi-unit agents, constraint violations will be minimal. As show in the next section, in practice, constraint violations can be much smaller than what is suggested by theoretical worst case described in Theorem 3.

However, Theorem 3 is silent about how much the values of localities’ objectives change in an integral group-stable matching when we round from a nearby fractional group-stable matching. In the proof of Theorem 3, we show that some probabilities that f is matched to ℓ (i.e. contracts) in the original fractional matching are rounded down to 0. This could potentially lead to a large loss in the value of the objective for such localities. Our next result shows that the loss in the values of localities’ objectives can also be bounded. Let us scale $v_{f,\ell}$ (i.e., the value of each family to each locality) to $\beta_{f,\ell}$ so that the scaled value of the objective of each locality with respect to coefficient $\beta_{f,\ell}$ and with respect to the dominating vector is 1. The reason for scaling the cardinal values of the objectives of localities to 1 is to guarantee comparability across localities. We are now ready to present our final result.

Theorem 4. *There exists a polynomial time algorithm that for any fractional group-stable matching finds an integral group-stable matching in which the total number of matched families does not decrease. Furthermore,*

- *the capacity of any service at any locality increases by at most α , and*

- value of the objective for each locality does not decrease by more than α ,

where $\alpha = \max_{f,\ell} (\beta_{f,\ell} + \sum_{s \in S} \alpha_{f,\ell}^s)$.

Proof. See Appendix B.3 □

Theorem 4 shows that it is possible to obtain an integral group-stable matching from any fractional group-stable that would neither significantly violate service capacity constraints nor cause a large decrease in value of the objective for any locality.

4.1 Iterative Rounding Algorithm

To prove Theorems 3 and 4, we introduce the Iterative Rounding (IR) algorithm which finds an integral group-stable matching from any fractional group-stable matching. This algorithm is similar to the one in Nguyen and Vohra (2018). The IR algorithm starts from a dominating vector (which may, of course, be fractional) and iteratively rounds this vector to a dominating integral solution. This rounding eventually produces an integral group-stable matching that may violate capacity constraints of some of the localities.

Specifically, the IR algorithm rounds the vector $(x_{f,\ell}^*)$ to an integer vector $(\bar{x}_{f,\ell})$. The integral vector $(\bar{x}_{f,\ell})$ corresponds to an integral matching. For $(\bar{x}_{f,\ell})$ to be group-stable with respect to the new constraints we require that $(\bar{x}_{f,\ell})$ satisfy properties stated in the following lemma.

Lemma 2. *Let $(x_{f,\ell,p}^*)$ be a (fractional) dominating vector of $A \cdot (x_{f,\ell,p}) \leq 1$. Then let $(\bar{x}_{f,\ell,p})$ be a non-negative integral vector satisfying the following properties:*

1. *for a family f and locality ℓ , if $x_{f,\ell,p}^* = 0$, then $\bar{x}_{f,\ell,p} = 0$,*
2. *for a family f if $\sum_{\ell,p} x_{f,\ell,p}^* = 1$ then $\sum_{\ell,p} \bar{x}_{f,\ell,p} = 1$.*

Finally, let $A \cdot (\bar{x}_{f,\ell,p}) = \kappa'$. Then $(\bar{x}_{f,\ell})$ obtained from $(\bar{x}_{f,\ell,p})$ is a group-stable matching with respect to κ' .

Proof. Property 1 ensures that the support of $(\bar{x}_{f,\ell,p})$ is contained within the support of $(x_{f,\ell,p}^*)$. Property 2 ensures that if a family is matched to a locality with probability 1 under $(x_{f,\ell,p}^*)$, it is also matched to that locality under $(\bar{x}_{f,\ell,p})$ with probability 1. This means that if a constraint binds under $(x_{f,\ell,p}^*)$, then it also binds under $(\bar{x}_{f,\ell,p})$. Both properties are necessary to make sure that the rounded solution $(\bar{x}_{f,\ell,p})$ continues to be a dominating vector with respect to the new capacity constraints. Then by Theorem 2, $(\bar{x}_{f,\ell})$ is a group-stable matching. \square

Let us provide some intuition behind the the IR algorithm before stating it formally. Throughout the IR algorithm, all integral components of vector x^* will remain fixed. At each non-terminal step, we solve an updated linear program that corresponds to the matching polytope and satisfies both properties stated in Lemma 2. The key observation here is that because of the sparsity of the constraint matrix, either (i) the solution of this updated linear program will have new integral components, or (ii) we can find a constraint that contains very few non-zero variables so even when we delete this constraint the new solution will not lead to a large violation of this constraint. If we are in Case (i), we fix the integral components of the solution. If we are in Case (ii), we delete the constraints and re-update the linear program. This iterative procedure continues until we round all fractional components of x^* either to 1 or to 0.

The IR algorithm is stated formally in Figure 1. In this algorithm the matrix A is the constraint matrix of the matching polytope. In the proof of Theorem 3 and Theorem 4 we describe the construction of matrix A in more detail.

5 Empirical illustration

We now illustrate our approach using real-world data from refugee resettlement. We also compare the performance of the our algorithm with a natural adaptation of the Deferred Acceptance (DA) algorithm.

Figure 1: ITERATIVE ROUNDING ALGORITHM (v, A)

v is a linear objective and A is the constraint matrix. We specify them later.

Step 0: Initiate $x^{opt} = x^*$.

Step 1: If x^{opt} is integral, stop and output x^{opt} ; otherwise continue either to Step 2a or 2b.

Step 2a: If any coordinate of x^{opt} is integral, fix the values of those coordinates, update the linear program as following and then move to *Step 3*.

Linear Program Update: Let B and \bar{B} be the set of columns of A that correspond to the non-integer- and integer-valued coordinates of x^{opt} respectively. Let A_B and $A_{\bar{B}}$ be the sub-matrix of A that consists of columns in A and the complement \bar{B} . For vector x , let x_B and $x_{\bar{B}}$ be the sub-vector of x that consists of all coordinates in B and \bar{B} . The updated linear program is:

$$\max\{v_B \cdot x_B : A_B \cdot x_B \leq 1 - A_{\bar{B}} \cdot x_{\bar{B}}^{opt}\}$$

In other words, we replace v by v_B , x by x_B , A by A_B , and 1 by $1 - A_{\bar{B}} \cdot x_{\bar{B}}^{opt}$.

Step 2b: If all coordinates of x^{opt} are fractional, we show that there is a row of A with a small row sum. (We specify it later.) Delete the corresponding constraint from the linear program. Update the linear program and move to *Step 3*.

Step 3: Solve the updated linear program to get an extreme point solution. Let this be the new x^{opt} and return to *Step 1*.

The algorithm terminates in a finite number of steps and outputs a 0-1 vector x^{OUT} (such that $v \cdot x^{OUT} \geq v \cdot x^*$).

5.1 Data

We used anonymized data on refugees and localities (known as “affiliates”) provided by the Hebrew Immigrant Aid Society (HIAS). The data cover resettlement of refugees without US ties for fiscal year 2017 (i.e., starting on October 1, 2016, and ending on September 30, 2017).

We dropped new localities and localities which only resettled a handful of refugees. We have

12 remaining localities. There are 200 families and 538 resettled refugees in our data. For every refugee family, we observe the number of adults, children, and seniors in that family. There are a total of 230 refugee children (under 18 years old), 304 refugee adults, and 4 refugee are seniors (over 65 years old) in our data.

5.2 Simulation setup

Currently, no data on refugees' preferences over localities are currently being collected. Therefore, we considered four different ways of generating refugee preferences.

1. **IDENTICAL**: in each simulation round, we generate a single preference list. We assign this preference list to each family. Therefore, in each simulation round, families' preferences are perfectly correlated.
2. **CORRELATED**: in each simulation round, we generate a single preference list. For each family, we pick several localities from this preference list and swap their order; we then assign the perturbed preference list to this family.⁶ Therefore, in each simulation round, families' preferences are imperfectly correlated.
3. **UNCORRELATED**: in each simulation round, generate a preference list for each family uniformly at random. Therefore, in each simulation round, families' preferences are uncorrelated.
4. **SIZE-BASED**: We generate a baseline preference list: rank all localities by their total capacity (sum of children capacity, adult capacity, and senior capacity) where larger localities are more preferred. Then, in each simulation round and for each family, we pick several localities from baseline preference list, swap their order and assign the perturbed preference list to this family. Therefore, families' preferences are correlated within and across simulations rounds.

⁶The number of pairs is equal to (the ceiling of) the number of localities that are compatible to the family divided by four.

Within a simulation round, the IDENTICAL produces perfect preference correlation, UNCORRELATED produces no preference correlation, and CORRELATED produces an intermediate level of preference correlation. However, in UNCORRELATED, CORRELATED, and IDENTICAL, preferences of refugees are independent across rounds. In SIZE-BASED, on the other hand, preferences are correlated within and across simulation rounds.

The value of family f at locality ℓ (i.e., $v_{f,\ell}$) is the estimated employment likelihood of that family at the locality. The employment likelihoods were estimated by [Trapp et al. \(2018\)](#) who applied several machine learning methods to eight years of HIAS resettlement data.

Actual resettlement numbers in each locality are typically close to the government-approved capacities. For our service constraints, we used the actual number of adults, children, and seniors resettled to each locality in fiscal year 2017. Therefore, each locality has a three-dimensional service constraint on the number of refugees it can resettle. We can consider these age-specific constraints as proxies for constraints on services provided in each locality (e.g., kindergarten and school places, job training, and senior care). Since each locality has three constraints, we have 36 constraints in total. The sparsity of the constraint matrix, (i.e., $\max_{f,\ell} (\sum_{s \in S} \alpha_{f,\ell}^s)$) is equal to 68%.

We compare the performance of our IR algorithm and of the (greedy) Deferred Acceptance (DA) algorithm after 100 simulation rounds. The DA algorithm works as follows. In each step, an unmatched family f proposes to the most preferred locality ℓ to which f has not yet proposed. Then, family f is added to locality ℓ 's choice set which already includes all families that are currently matched with the locality. Locality ℓ then chooses a subset of families from its choice set that maximizes its objective function. If a family is not chosen, it is rejected. The DA algorithm terminates when there is no unmatched family or every family has proposed to all localities on its list.

The simulations were coded in Matlab and run on a desktop computer (CPU: 3.6 GHz AMD Ryzen 5 3600 Processor with 6 physical cores; RAM: 16 GB). A fractional matching

was typically found in about 2 hours in each simulation round. The IR algorithm and the DA algorithm ran in a matter of seconds. Thus, the heaviest computation burden is finding the fractional stable matching via Scarf’s Lemma. Finding a dominating solution in general is known to be PPAD-hard (Kintali, 2008). However, as shown in our simulations, it is practical for the refugee resettlement problem. In a related problem of matching with couples the Scarf algorithm has been shown to computationally outperform other popular algorithms (Biró et al., 2016).

5.3 Results

Tables 2 and 3 summarize the performance of the IR and the DA algorithms under all preference specifications.

Table 2 shows that the IR algorithm outperforms the DA algorithm with respect to the number of matched families, the number of matched refugees, and total employment. Interestingly, the largest performance gap is for the case of IDENTICAL preferences where the DA algorithm is equivalent to a serial dictatorship and does not produce any blocking pairs.

Table 3 summarizes constraint violations under the IR algorithm. Across all four preference specifications, the maximum capacity constraint violation was 44%: a constraint (on the number of children) of capacity 9 was relaxed by 4 units. This maximal capacity violation is still substantially below the sparsity of the constraint matrix. However, the average capacity constraint violations are substantially smaller than the largest violations. For example, the average maximum capacity constraint violations across all simulation rounds is 13.9% while the average violation of violated capacity constraint is just 7.3%. If we consider all constraints (by assigning a zero percent violation to any constraint that is not violated), then the average capacity constraint violation is a meager 0.9%. On average, only 5 constraints are violated while 16 constraints are never violated.

Figures 2 and 3 illustrate the violation of constraints on the number of children and

	UNCORRELATED		CORRELATED		IDENTICAL		SIZE-BASED	
	DA	SCARF+IR	DA	SCARF+IR	DA	SCARF+IR	DA	SCARF+IR
Matched families	189 (1.4)	193 (1.4)	185 (2.2)	190 (1.8)	176 (2.5)	184 (3.0)	189 (1.3)	193 (1.5)
Matched refugees	487 (8.4)	509 (8.5)	467 (12.1)	497 (11.4)	427 (12.1)	470 (10.8)	468 (7.2)	509 (9.4)
Total employment	93.00 (1.9)	93.97 (1.7)	93.45 (1.9)	94.84 (1.7)	92.57 (1.9)	96.16 (1.7)	90.58 (1.4)	92.29 (1.4)
Blocking pairs	281 (101.3)	0	210 (93.4)	0	0	0	224 (65.1)	0

Table 2: Summary of simulation results: size of matchings, total employment, and stability. Number of simulation rounds for each preference specification: 100. Standard deviations in parentheses.

	UNCORRELATED	CORRELATED	IDENTICAL	SIZE-BASED	AVERAGE
Ave. cap. violation across all constraints	1.0%	1.2%	0.8%	0.7%	0.9 %
Ave. cap. violation across violated constraints	7.6%	9.4%	6.4%	5.9%	7.3%
Ave. maximum cap. violation	12.7 %	20.0 %	12.9%	9.8%	13.9 %
Maximum cap. violation	33.3%	33.3 %	44.4%	44.4%	38.9%
Cap. of constraint with maximum cap. violation	9	9	9	9	9
Number of constraints that are never violated	17	17	14	16	16
Ave. number of violated constraints	5	5	4	7	5
Number of violations of the most violated constraint	64	64	52	67	62

Table 3: Summary of constraint violations. Number of simulation rounds for each preference specification: 100. AVERAGE averages over all 400 simulation rounds. Number of constraints: 36.

1. For every simulation round, sum all capacity violations (assigning a zero percent violation to any constraint that is not violated). Divide by 3600.
2. For every simulation round, sum all capacity violations of violated constraints. Divide by 3600.
3. Find the largest capacity violation across all constraints and across all simulation rounds.
4. For every simulation round, find the largest capacity violation across all constraints. Divide by 100.
5. Find the capacity of the constraint with the largest capacity constraint violation across all simulation rounds.
6. Find the size of the set constraints are not violated in any simulation round.
7. For every simulation round, find the number of violated constraints. Sum and divide by 100.
8. For every constraint, count the number of times the constraint has been violated across all simulation rounds. Take the maximum over all constraints.




Figure 2: Violation of the capacity constraint on the number of children: Average capacity under IR algorithm vs. actual capacity across all localities. Circles indicate localities where the constraint was violated (on average across all simulation rounds).

number of adults across all localities (the constraint on seniors is never violated). The constraint on the number of children is violated (on average) in four localities while the constraint on the number of adults is violated in one locality.

Overall, the extent of constraint violation produced by the our algorithm makes our algorithm potentially relevant for refugee resettlement. In practice, American localities are permitted to exceed their government approved annual quota by at most 10% (Trapp et al., 2018). While in some instances of the simulations our algorithm produces greater than 10% violation for some small capacity constraints, the results are encouraging for practical implementation.

6 Conclusion

This paper analyzed fractional group-stable matchings in a model of many-to-one matching with multidimensional knapsack constraints. Using Scarf’s Lemma, we developed an algorithm that finds an integral group-stable matching by minimally relaxing of these knapsack constraints. The constraint violation and the value loss are proportional to the relative “size”




Figure 3: Violation of the capacity constraint on the number of adults: Average capacity under IR algorithm vs. actual capacity across all localities. Circles indicate localities where the constraint was violated (on average across all simulation rounds).

of the “largest” agent ensuring good performance in a variety of important applications.

Our simulation results based on refugee resettlement data suggest that the error bounds in practice are much smaller than the worst case provided by the theory. Hence, our algorithm offers a new practical solution, which is more stable and can be more efficient than the outcomes generated by the (greedy) Deferred Acceptance algorithm.

There are several directions for further work. First, one could introduce minimum quotas or other types of complex distributional constraints. Second, one could consider nonlinear objective functions for multi-unit agents or allow unit agents to have preferences over other unit agents. Third, one could study strategic properties of our solution in large markets.

APPENDIX

A Complementary Slackness Theorem

Lemma A.1. *Let (P) and (D) be the primal and dual linear programming problems, respectively.*

$$\max\{c^T x : Ax \leq b, x \geq 0\} \quad (P)$$

$$\min\{b^T x : A^T y \geq c, y \geq 0\} \quad (D)$$

Then x^ and y^* are optimal of (P) and (D) , respectively, if and only if the following conditions are satisfied.*

$$\left(b_i - \sum_j a_{ij}x_j\right) y_i = 0 \text{ and } \left(\sum_i a_{ij}y_i - c_j\right) x_j = 0.$$

B Proofs

B.1 Proofs of Claims 1 and 2

Here we prove the two claims needed for Theorem 1. The first claim says that if the primal variable is slack, then the corresponding dual constraint must be binding when ϵ goes to 0.

Proof of Claim 1. Consider two cases:

Case 1: $x_{f,\ell}^* = 1$, then $\bar{q}_f + \sum_{s \in S} \alpha_{f,\ell}^s \cdot \bar{p}_\ell^s = v_{f,\ell}$.

Case 2: $x_{f,\ell}^* < 1$, then $\bar{q}_f = 0$. For the sake of contradiction, suppose $\sum_{s \in S} \alpha_{f,\ell}^s \cdot \bar{p}_\ell^s < v_{f,\ell} - \epsilon \sum_{s \in S} \alpha_{f,\ell}^s$. We look at the price vector \hat{p} in which each constraint price is ϵ higher than the minimum price \bar{p}_ℓ^s . Then clearly \hat{p}_ℓ is a feasible price vector for locality ℓ and family f since $\sum_{s \in S} \alpha_{f,\ell}^s \cdot \hat{p}_\ell^s < v_{f,\ell}$. It implies that there exists a column (f, ℓ, \hat{p}) in \mathcal{Q} corresponding to that price. Now for any service $s \in S$, the row (ℓ, s) prefers column (f, ℓ, \hat{p}) to the column

containing the minimum price \bar{p}_ℓ^s . Note that the column (f, ℓ, \hat{p}) cannot be dominated by any row corresponding to a family since $x_{f,\ell}^* < 1$. Therefore, column (f, ℓ, \hat{p}) is not dominated by x^* , a contradiction. \square

The second claim says that if the primal constraint is slack, then the corresponding dual variable must be zero when ϵ goes to 0.

Proof of Claim 2. The latter statement is trivial by the definition of \bar{q}_f , we just need to prove the former statement. Suppose it is not true. We look at the price whose each constraint price for any service $s' \neq s$ is ϵ higher than the minimum price and the constraint price for service s is lower by the total gap between other constraints. Then it is easy to see that there exists a column corresponding to that price, and that column is not dominated, contradiction. \square

B.2 Proof of Theorem 3

Proof of Theorem 3. We first show how to apply IR algorithm. Let the linear objective v capture the total number of families assigned in a matching. Note that we can also weigh these families according to their size.

Recall that $x_{f,\ell}$ the sum of all the contracts between family f and locality ℓ over all possible prices, $x_{f,\ell} = \sum_p x_{f,\ell,p}$. Moreover, recall from the proof of Theorem 2 that equations (2) and (3) are respectively equivalent to

$$\sum_{f \in F} \alpha_{f,\ell}^s \cdot x_{f,\ell} \leq 1 \quad \text{for every } \ell \in L \text{ and } s \in S, \quad (\text{B.1})$$

$$\sum_{\ell \in L} x_{f,\ell} \leq 1 \quad \text{for every } f \in F. \quad (\text{B.2})$$

We apply the (IR) algorithm to the system (B.1)–(B.2) with some utility function u where $u \geq 0$, for example, u is the total number of matched families $u = \sum_{f,\ell} x_{f,\ell}$. Let A be the constraint matrix associated with locality constraint (B.1). Then the constraint can be

expressed as $Ax \leq 1$. Define approximate constraint of (B.1) as

$$Ax \leq \kappa', \tag{B.3}$$

where $\kappa' \leq 1 + \alpha$ and $\alpha = \max_f \sum_{s \in S} \alpha_{f,\ell}^s$. The IR algorithm takes as input an extreme point of (B.1)–(B.2), let us call it $(x_{f,\ell}^*)$. It then rounds $(x_{f,\ell}^*)$ into a 0-1 vector $(\bar{x}_{f,\ell})$ that satisfies (B.2) and (B.3). Note that at an extreme point there does not exist any pair of (f, ℓ) such that there are two contracts $x_{(f,\ell,p)} > 0$ and $x_{(f,\ell,p')} > 0$. It means that for each $x_{f,\ell} > 0$, there is a unique price p such that $x_{(f,\ell,p)} > 0$. So after rounding, we can use the same price to set the value of $\bar{x}_{f,\ell,p} = x_{f,\ell,p}^*$. It is easy to see that $(\bar{x}_{f,\ell,p})$ satisfies Lemma 2. Note that to make sure that the second property in Lemma 2 holds when we apply the IR algorithm we require that if initially a constraint was binding with respect to $(x_{f,\ell}^*)$ (resp. $(x_{f,\ell,p}^*)$) then that constraint must be binding with respect to $(\bar{x}_{f,\ell})$ (resp. $(\bar{x}_{f,\ell,p})$). Therefore when updating the linear program in Step (2a) of Fig. 1 we ensure that all equalities are preserved.

Beginning with $(x_{f,\ell}^*)$, the algorithm executes Step (2a). Remove from the linear program all variables $x_{f,\ell}$ that are integral and adjust the right hand sides of all the constraints in the linear program accordingly. Thus, a variable set to 0 or 1 by $(x_{f,\ell}^*)$ is fixed at 0 or 1 respectively in all subsequent iterations. In the remaining system, pick a non-negative extreme point that optimizes the vector u and repeat. At some iteration we may in an extreme point in which each variable is non-integer, call it $(x_{f,\ell}^*)$. At this point we must execute Step (2b), identifying a constraint to be removed. We show that there always exists a constraint such that

$$\sum_f \alpha_{f,\ell}^s \cdot \lceil x_{f,\ell}^* \rceil \leq 1 + \alpha.$$

To prove that, assume for the sake of contradiction that

$$\sum_f \alpha_{f,\ell}^s \cdot \lceil x_{f,\ell}^* \rceil > 1 + \alpha. \quad (\text{B.4})$$

We use (B.4) to contradict the fact that $(x_{f,\ell}^*)$ is an extreme point. We will use the following property: The number of non-zero variables in an extreme point is equal to the number of linearly independent and binding constraints.

Given the extreme point $(x_{f,\ell}^*)$, we credit each non-zero variable with a single token. We then redistribute the tokens to binding, linearly independent constraints in the following way:

- credit $x_{f,\ell}^*$ token to the constraint corresponding to (B.2); and
- credit $\alpha_{f,\ell}^s \cdot \frac{1-x_{f,\ell}^*}{\alpha}$ to each constraint of ℓ corresponding to service s .

If the constraint corresponding to (B.2) binds, then the number of tokens this constraint credited with is $\sum x_{f,\ell}^* = 1$. For a binding constraint corresponding to locality ℓ and service s , we have

$$\sum_f \alpha_{f,\ell}^s x_{f,\ell}^* = 1.$$

The total quantity of tokens credited to this constraint is

$$\sum_f \alpha_{f,\ell}^s \frac{1-x_{f,\ell}^*}{\alpha} > \frac{1}{\alpha}(1 + \alpha - 1) = 1.$$

Thus, any binding constraint in (B.1) is credited with more than 1 token. It implies that the number of non-zero variables is larger than the number of linearly independent binding constraints.

To conclude, since $(\bar{x}_{f,\ell,p})$ satisfies Lemma 2, $(\bar{x}_{f,\ell})$ is an integral group-stable matching.

During the IR algorithm we only drop constraints, hence the objective function $v \cdot x$ does

not decrease. This shows that the total number of families assigned by integral group-stable matching is at least as large as the number of families assigned in the fractional group-stable matching. \square

B.3 Proof of Theorem 4

Proof of Theorem 4. To show that value of the objective for each locality does not decrease too much after rounding, we add the following constraint for each locality.

$$\sum_{f \in F} \beta_{f,\ell} \cdot x_{f,\ell} \geq 1. \tag{B.5}$$

Then we apply the (IR) algorithm to the system of inequalities (B.1), (B.2), (B.5) with the utility function u . The only difference in this case compared to the proof of Theorem 3 is that when executing Step (2b) the constraint to be removed can be either

1. A constraint satisfying $\sum_f \alpha_{f,\ell}^s \cdot \lceil x_{f,\ell}^* \rceil \leq 1 + \alpha$ (like in previous case).
2. Or a constraint satisfying $\sum_f \beta_{f,\ell} \cdot x_{f,\ell}^* \leq \alpha$.

Intuitively, the Constraint 1 ensures that the rounded solution does not violate the constraint by no more than α , while Constraint 2 ensures that the total utility of each locality does not decrease by more than α .

Now, we will prove that there always exists a constraint satisfying either one of the two conditions above. Again, for the sake of contradiction, assume that

$$\sum_f \alpha_{f,\ell}^s \cdot \lceil x_{f,\ell}^* \rceil > 1 + \alpha, \text{ and}$$

$$\sum_f \beta_{f,\ell} \cdot x_{f,\ell}^* > \alpha.$$

We will use these two conditions to contradict the fact that $(x_{f,\ell}^*)$ is an extreme point. We make the following modification to the token distribution: Given the extreme point $(x_{f,\ell}^*)$, we credit each non-zero variable a single token. We then redistribute the tokens to binding, linearly independent constraints in a particular way:

- credit $\alpha_{f,\ell}^s \cdot \frac{1-x_{f,\ell}^*}{\alpha}$ to each constraint ℓ corresponding to service s ;
- credit $\beta_{f,\ell} \cdot \frac{x_{f,\ell}^*}{\alpha}$ to the constraint corresponding to (B.5); and
- credit $\frac{\beta_{f,\ell}}{\alpha}(1-x_{f,\ell}^*) + (1-\frac{\beta_{f,\ell}}{\alpha})x_{f,\ell}^*$ to (B.2).

If the constraint corresponding to (B.1) binds, then $\sum_f \alpha_{f,\ell}^s x_{f,\ell}^* = 1$. Hence, the total quantity of tokens credited to this constraint is

$$\sum_f \alpha_{f,\ell}^s (1-x_{f,\ell}^*) = \frac{1}{\alpha} \left(\sum_f \alpha_{f,\ell}^s - \sum_f \alpha_{f,\ell}^s x_{f,\ell}^* \right) > \frac{1}{\alpha} (1 + \alpha - 1) = 1.$$

The total quantity of tokens obtained by constraint corresponding to (B.5) is

$$\sum_f \beta_{f,\ell} \cdot \frac{x_{f,\ell}^*}{\alpha} = \frac{1}{\alpha} \sum_f \beta_{f,\ell} \cdot x_{f,\ell}^* > \frac{1}{\alpha} \cdot \alpha = 1.$$

Finally, if the constraint corresponding to (B.2) binds, then $\sum_{\ell \in L} x_{f,\ell}^* = 1$. If each constraint is credited $x_{f,\ell}^*$, then clearly the total quantity of tokens that constraint obtained is 1. If each constrained is credited $1-x_{f,\ell}^*$, then the total quantity of tokens that constraint obtained is $\sum_{\ell} (1-x_{f,\ell}^*) = \sum_{\ell} 1 - \sum_{\ell} x_{f,\ell}^* \geq 2 - 1 = 1$ since there are always at least two localities in fractional case. Note that $\frac{\beta_{f,\ell}}{\alpha}(1-x_{f,\ell}^*) + (1-\frac{\beta_{f,\ell}}{\alpha})x_{f,\ell}^*$ is a convex combination of $(1-x_{f,\ell}^*)$ and $x_{f,\ell}^*$. Therefore, the total quantity of tokens this constraint obtained is at least 1.

Hence, any binding constraint is credited with more than 1 token. This implies that the number of non-zero variables is larger than the number of linearly independent binding constraints, a contradiction.

Because we remove a utility constraint only when $\sum_f \beta_{f,\ell} \cdot x_{f,\ell}^* \leq \alpha$, and in the worst case, $x_{f,\ell}^*$ are rounded down to 0, we throw away at most α when rounding. This implies that the value of the objective for each locality does not decrease by more than α .

Similar to the argument in the proof of Theorem 3, we choose the linear objective v to capture the total (weighted) number of families assigned. During the IR algorithm we only drop constraints, hence the objective function $v \cdot x$ does not decrease. This shows that the total number of families assigned by integral group-stable matching is at least large as the number of families assigned in the fractional group-stable matching. \square

References

- Abdulkadiroğlu, A. and T. Sönmez (2003). School choice: A mechanism design approach. *American Economic Review* 93(3), 729–747.
- Aharoni, R. and T. Fleiner (2003). On a lemma of Scarf. *Journal of Combinatorial Theory, Series B* 87(1), 72–80.
- Aharoni, R. and R. Holzman (1998). Fractional kernels in digraphs. *Journal of Combinatorial Theory, Series B* 73(1), 1–6.
- Ashlagi, I., A. Saberi, and A. Sharneli (2020). Assignment mechanisms under distributional constraints. *Operations Research* 68(2), 467–479.
- Ashlagi, I. and P. Shi (2016). Optimal allocation without money: An engineering approach. *Management Science* 62(4), 1078–1097.
- Aygiin, O. and I. Bó (2019, July). College admission with multidimensional privileges: The Brazilian affirmative action case. WZB Berlin Social Science Center Working Paper.
- Aygiin, O. and B. Turhan (2017). Large-scale affirmative action in school choice: Admissions to IITs in India. *American Economic Review* 107(5), 210–13.
- Azevedo, E. M. and J. D. Leshno (2016). A supply and demand framework for two-sided matching markets. *Journal of Political Economy* 124(5), 1235–1268.
- Bansak, K., J. Ferwerda, J. Hainmueller, A. Dillon, D. Hangartner, D. Lawrence, and J. Weinstein (2018). Improving refugee integration through data-driven algorithmic assignment. *Science* 359(6373), 325–329.
- Biró, P. and T. Fleiner (2016). Fractional solutions for capacitated NTU-games, with applications to stable matchings. *Discrete Optimization* 22, 241–254.

- Biró, P., T. Fleiner, and R. W. Irving (2016). Matching couples with Scarf’s algorithm. *Annals of Mathematics and Artificial Intelligence* 77(3-4), 303–316.
- Biró, P., T. Fleiner, R. W. Irving, and D. F. Manlove (2010). The college admissions problem with lower and common quotas. *Theoretical Computer Science* 411(34), 3136 – 3153.
- Biró, P. and J. Gudmundsson (2020, forthcoming). Complexity of finding Pareto-efficient allocations of highest welfare. *European Journal of Operations Research*.
- Biró, P. and E. McDermid (2014). Matching with sizes (or scheduling with processing set restrictions). *Discrete Applied Mathematics* 164, 61–67.
- Bronfman, S., N. Alon, A. Hassidim, and A. Romm (2018). Redesigning the Israeli medical internship match. *ACM Transactions on Economics and Computation (TEAC)* 6(3-4), 21.
- Che, Y.-K., J. Kim, and F. Kojima (2018). Stable matching in large economies. *Econometrica*.
- Che, Y.-K. and O. Tercieux (2018). Payoff equivalence of efficient mechanisms in large matching markets. *Theoretical Economics* 13(1), 239–271.
- Correa, J., R. Epstein, J. Escobar, I. Rios, B. Bahamondes, C. Bonet, N. Epstein, N. Aramayo, M. Castillo, A. Cristi, et al. (2019). School choice in Chile. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pp. 325–343.
- Dean, B. C., M. X. Goemans, and N. Immorlica (2006). The unsplitable stable marriage problem. In *Fourth IFIP International Conference on Theoretical Computer Science-TCS 2006*, pp. 65–75. Springer.
- Delacrétaz, D. (2019, January). Stability in matching markets with sizes.
- Delacrétaz, D., S. D. Kominers, and A. Teytelboym (2019, December). Matching mechanisms for refugee resettlement. Oxford University Working Paper.
- Echenique, F. and M. B. Yenmez (2015). How to control controlled school choice. *The American Economic Review* 105(8), 2679–2694.
- Ehlers, L., I. E. Hafalir, M. B. Yenmez, and M. A. Yildirim (2014). School choice with controlled choice constraints: Hard bounds versus soft bounds. *Journal of Economic Theory* 153, 648–683.
- Erdil, A. and T. Kumano (2012). Prioritizing diversity in school choice. Working paper, Washington University.
- Fragiadakis, D., A. Iwasaki, P. Troyan, S. Ueda, and M. Yokoo (2016). Strategyproof matching with minimum quotas. *ACM Transactions on Economics and Computation* 4(1), 6.
- Gale, D. and L. S. Shapley (1962). College admissions and the stability of marriage. *American Mathematical Monthly* 69(1), 9–15.

- Gonczarowski, Y. A., N. Nisan, L. Kovalio, and A. Romm (2019). Matching for the Israeli “mechinot” gap-year programs: Handling rich diversity requirements. In *Proceedings of the 2019 ACM Conference on Economics and Computation, EC ’19*, New York, NY, USA, pp. 321–321. ACM.
- Hassidim, A., A. Romm, and R. I. Shorrer (2017). Redesigning the Israeli psychology master’s match. *American Economic Review* 107(5), 205–09.
- Hatfield, J. W. and F. Kojima (2008). Matching with contracts: Comment. *American Economic Review* 98(3), 1189–94.
- Hatfield, J. W. and P. Milgrom (2005). Matching with contracts. *American Economic Review* 95(4), 913–935.
- Huang, C.-C. (2010). Classified stable matching. In *Proceedings of the twenty-first annual ACM-SIAM symposium on Discrete Algorithms*, pp. 1235–1253. SIAM.
- Irving, R. W., P. Leather, and D. Gusfield (1987). An efficient algorithm for the optimal stable marriage. *Journal of the ACM (JACM)* 34(3), 532–543.
- Jagadeesan, R. (2017). Complementary inputs and the existence of stable outcomes in large trading networks. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pp. 265–265. ACM.
- Kamada, Y. and F. Kojima (2012). Stability and strategy-proofness for matching with constraints: A problem in the Japanese medical match and its solution. *American Economic Review* 102(3), 366–370.
- Kamada, Y. and F. Kojima (2015). Efficient matching under distributional constraints: Theory and applications. *American Economic Review* 105(1), 67–99.
- Kamada, Y. and F. Kojima (2019, January). Fair matching under constraints: Theory and applications. Working paper.
- Kelso, A. S. and V. P. Crawford (1982). Job matching, coalition formation, and gross substitutes. *Econometrica* 50(6), 1483–1504.
- Kintali, S. (2008). Scarf is ppad-complete. *arXiv preprint arXiv:0812.1601*.
- Király, T. and J. Pap (2008). Kernels, stable matchings, and Scarf’s Lemma. Technical Report TR-2008-13, Egerváry Research Group.
- Klaus, B. and F. Klijn (2005). Stable matchings and preferences of couples. *Journal of Economic Theory* 121(1), 75 – 106.
- Milgrom, P. (2017, May). *Discovering Prices: Auction Design in Markets with Complex Constraints*. Columbia University Press.
- Milgrom, P. and I. Segal (2020). Clock auctions and radio spectrum reallocation. *Journal of Political Economy* 128(1), 1–31.

- Nguyen, T., A. Peivandi, and R. Vohra (2016). Assignment problems with complementarities. *Journal of Economic Theory* 165, 209–241.
- Nguyen, T. and R. Vohra (2018). Near-feasible stable matchings with couples. *American Economic Review* 108(11), 3154–69.
- Nguyen, T. and R. Vohra (2019). Stable matching with proportionality constraints. *Operations Research* 67(6), 1503–1519.
- Noda, S. (2018). Large matchings in large markets with flexible supply. *Available at SSRN 3215670*.
- Roth, A. E. (1984). The evolution of the labor market for medical interns and residents: A case study in game theory. *Journal of Political Economy* 92(6), 991–1016.
- Roth, A. E. (1991). A natural experiment in the organization of entry-level labor markets: Regional markets for new physicians and surgeons in the United Kingdom. *American Economic Review* 81, 415–440.
- Roth, A. E. (2002). The economist as engineer: Game theory, experimentation, and computation as tools for design economics. *Econometrica* 70(4), 1341–1378.
- Roth, A. E., U. G. Rothblum, and J. H. Vande Vate (1993). Stable matchings, optimal assignments, and linear programming. *Mathematics of Operations Research* 18(4), 803–828.
- Roth, A. E. and X. Xing (1994). Jumping the gun: Imperfections and institutions related to the timing of market transactions. *American Economic Review* 84, 992–1044.
- Roth, B. and R. I. Shorrer (2018). Making it safe to use centralized marketplaces: Dominant individual rationality and applications to market design. *Available at SSRN 3073027*.
- Scarf, H. E. (1967). The core of an n -person game. *Econometrica* 35(1), 50–69.
- Sethuraman, J., C.-P. Teo, and L. Qian (2006). Many-to-one stable matching: geometry and fairness. *Mathematics of Operations Research* 31(3), 581–596.
- Shi, P. (2015). Guiding school-choice reform through novel applications of operations research. *Interfaces* 45(2), 117–132.
- Sönmez, T. and M. B. Yenmez (2019a). Affirmative action with overlapping reserves. Technical report, Boston College Department of Economics.
- Sönmez, T. and M. B. Yenmez (2019b). Constitutional Implementation of Vertical and Horizontal Reservations in India: A Unified Mechanism for Civil Service Allocation and College Admissions. Working paper, Boston College Department of Economics.
- Tan, J. J. (1991). A necessary and sufficient condition for the existence of a complete stable matching. *Journal of Algorithms* 12(1), 154–178.

- Teo, C.-P. and J. Sethuraman (1998). The geometry of fractional stable matchings and its applications. *Mathematics of Operations Research* 23(4), 874–891.
- Trapp, A. C., A. Teytelboym, A. Martinello, T. Andersson, and N. Ahani (2018, October). Placement optimization in refugee resettlement. Working Paper 2018:23, Lund University, Department of Economics.
- Veski, A., P. Biró, K. Pöder, and T. Lauri (2017). Efficiency and fair access in kindergarten allocation policy design. *The Journal of Mechanism and Institution Design* 2(1), 57–104.
- Wu, X. (2018). Core of convex matching games: A Scarf’s Lemma approach. Working paper, Columbia University, Department of Economics.