

Matching Mechanisms for Refugee Resettlement*

David Delacrétaz[†]
University of Oxford

Scott Duke Kominers[‡]
Harvard Business School

Alexander Teytelboym[§]
University of Oxford

September 21, 2020

Abstract

Tens of thousands of refugees are permanently resettled from refugee camps to hosting countries every year. In the past, placement of refugees was essentially ad hoc, but more recently resettlement agencies have been trying to place refugees systematically in order to improve their outcomes. Yet, even at present, refugee resettlement processes account for neither the priorities of hosting communities nor the preferences of refugees themselves. Building on models from two-sided matching theory, we introduce a new framework for *matching with multidimensional knapsack constraints* that models the (possibly multidimensional) sizes of families, as well as the capacities of localities. We propose four refugee resettlement mechanisms and two solution concepts that can be used in refugee resettlement matching under various institutional and informational constraints. Our theoretical results and simulations using refugee resettlement data suggest that preference-based matching mechanisms can improve match efficiency, respect priorities of localities, and incentivize refugees to report where they would prefer to settle.

*First version: November 8, 2016 under the title “Refugee Resettlement.” The authors appreciate the helpful comments of Tommy Andersson, Ramnik Arora, Georgy Artermov, Haris Aziz, Ivan Balbuzanov, Péter Biró, Vincent Crawford, Sarah Glatte, Jens Gudmundsson, Guillaume Haeringer, Cameron Hepburn, Will Jones, Fuhito Kojima, Bettina Klaus, Simon Loertscher, Mike Mitchell, Karen Monken, Alex Nichifor, Assaf Romm, Alvin Roth, Yang Song, Tayfun Sönmez, Bassel Tarbush, Andrew Trapp, William Thomson, Utku Ünver, and Alex Westkamp, as well as four referees, Pol Antràs, Andrei Shleifer, and numerous conference and seminar participants. We are very grateful to Hai Nguyen for terrific research assistance.

[†]Email: david.delacretaz@economics.ox.ac.uk. Much of this work was completed when Delacrétaz was at the University of Melbourne. Delacrétaz is grateful for the support of the Faculty of Business and Economics, the Department of Economics, and the Centre for Market Design at the University of Melbourne, as well as the Australian Research Council grant DP160101350.

[‡]Email: skominers@hbs.edu. Kominers is grateful for the support of National Science Foundation grants CCF-1216095 and SES-1459912, the Harvard Milton Fund, the Washington Center for Equitable Growth, the Ng Fund and the Mathematics in Economics Research Fund of the Harvard Center of Mathematical Sciences and Applications, and the Human Capital and Economic Opportunity Working Group (HCEO) sponsored by the Institute for New Economic Thinking (INET).

[§]Email: alexander.teytelboym@economics.ox.ac.uk. Teytelboym is grateful for the support of the Skoll Centre for Social Entrepreneurship at the Saïd Business School, as well as for the generous fellowship and hospitality of the EU Centre for Shared Complex Challenges at the University of Melbourne. This work was supported by the Economic and Social Research Council grant number ES/R007470/1.

1 Introduction

At the end of 2019, 79.5 million people were displaced by conflict around the world—the highest level ever recorded (UNHCR, 2020). Over 20 million of these forcibly displaced people are deemed to be *refugees* under the mandate of the United Nations High Commissioner for Refugees (UNHCR). The UNHCR estimates that, in 2020, around 1.44 million refugees will not be able return to their home countries safely in the future (UNHCR, 2019a). The UNHCR deems these refugees eligible for *resettlement* in states that agree to give them permanent residence and a route to citizenship. Refugees eligible for resettlement are some of the most vulnerable refugees in the world, including children, survivors of torture and persecution, as well as women and girls at risk of violence (UNHCR, 2019b). While the number of refugees eligible for resettlement has roughly doubled since 2012, the number of resettled refugees has been falling in recent years. Resettlement places are provided by countries voluntarily—and the largest hosts are the United States, Canada, the United Kingdom, France, Sweden, and Australia. Since the Second World War, the U.S. has admitted the majority of resettled refugees. For example, between 2012 and 2018, the U.S. admitted an average of 46,000 refugees every year.

Yet little attention has been paid to the process that determines where in the hosting country refugees are resettled. Most countries have historically treated refugee resettlement as a purely administrative issue—and as such, they have not developed systematic and transparent resettlement procedures. There is, however, ample empirical evidence that the local communities (*localities*) where refugees are initially resettled matter a great deal for refugees’ education, job prospects, and earnings (Åslund and Rooth, 2007; Åslund and Fredriksson, 2009; Åslund et al., 2010, 2011; Damm, 2014; Feywerda and Gest, 2016; Bansak et al., 2018; Martén et al., 2019). In particular, the initial match matters for refugees’ lifetime outcomes because most refugees do not move from their initial resettlement localities for many years.

As of May 2018, HIAS (founded as the Hebrew Immigrant Aid Society), a resettlement agency operating in the U.S., has been systematically matching refugees according to their likelihood of gaining employment while taking various integration constraints into account (Trapp et al., 2020). Such systematic matching has the potential to increase the short-term employment of resettled refugees from 30 percent to over 40 percent while ensuring that all locality-level constraints are met (Bansak et al., 2018; Trapp et al., 2020).

HIAS’s pioneering matching software, Annie™ MOORE, is based on a model of matching with multidimensional knapsack constraints that we introduced in our original working paper (Delacrétaz et al., 2016; Trapp et al., 2020).¹ However, Annie™ MOORE at present neither

¹In this case our model can be analyzed as an integer program called the *multiple multidimensional*

accounts for refugees’ preferences over localities nor the priorities of the localities themselves.

Refugees’ preferences matter because refugees have private information about their own skills and abilities, which can affect the refugee-locality match quality—and which cannot be directly observed by government authorities. Meanwhile, respecting priorities and hosting capacities of localities can improve integration-relevant outcomes of refugees, ensure the best use of local resources, and encourage localities to continue participating in resettlement by building community support.

In this paper, we consider how refugees’ preferences and localities’ priorities can be incorporated into refugee resettlement processes, such as the one used by HIAS. We introduce and analyze several matching market design approaches that balance competing objectives of refugee welfare, incentives, and respect for localities’ priorities.

Our analysis draws in part upon classic matching models from contexts such as the matching of students to (public) schools. But in school choice contexts, any given student takes up exactly one school seat.² In refugee resettlement, by contrast, families must be kept together. Therefore, families have different *sizes*: larger families (e.g., a couple with four children) take up more places than smaller ones (e.g., a single individual).

Because localities in practice have inflexible total quotas on the number of refugees they are willing to admit, family sizes render most standard matching mechanisms (e.g., for the allocation of school seats, houses, or other objects) insufficient for refugee resettlement. Additionally, localities often have further requirements that constrain the allocation, such as a maximum number of single-parent families, refugees speaking a given language, or school-age children that they can accommodate. Thus, in order to deal with heterogeneous family sizes and various additional requirements, we allow for (possibly multidimensional) *knapsack constraints* that limit the central authority’s ability to allocate refugees to localities simply on the basis of numbers of individual refugees.³ In this framework, we propose matching mechanisms that account for preferences and priorities while respecting the knapsack constraints.

Our theoretical contribution. We consider a general model with multidimensional knapsack constraints, in which each family has a (possibly multidimensional) size and each locality

knapsack problem (Delacrétaz et al., 2016; Trapp et al., 2020).

²Students have heterogeneous preferences over schools and schools have priorities over students (having a sibling or living in the neighborhood typically gives students a higher priority). The social planner’s objective is to elicit truthful preferences over schools from students (schools are assumed to be non-strategic and school seats are treated as objects) and to deliver a non-wasteful matching of students to schools in which no student envies another student’s seat.

³The version of our model in which the only constraint in each locality is a maximum number of refugees is sufficient for HIAS’s current matching problem; however, the additional generality of our model may prove useful for future development of refugee resettlement matching (see Section 3.1 for a discussion).

has a (possibly multidimensional) capacity. A group of families can only be matched to a locality if, for every dimension, the total size of the families does not exceed the locality’s capacity. The implementation of our model by HIAS via Annie™ MOORE constitutes a special case of our model in which there is only one dimension: the total number of refugees (Trapp et al., 2020). All our results and solution concepts are novel even when there is only one dimension. Indeed, with the exception of Theorems 2 and 4 and Proposition 8, our results are not even affected by the number of dimensions.

We start with the case in which the resettlement agency focuses only on the preferences of refugee families. We show that the Knapsack Top Trading Cycles (KTTC) algorithm—a slight modification of the classical Top Trading Cycles (TTC) algorithm of Shapley and Scarf (1974)—allows us to incorporate knapsack constraints and obtain a Pareto-efficient mechanism in which refugee families do not have any incentive to misreport their preferences (Proposition 1). In practice, however, resettlement agencies already have existing allocation processes so we consider how to incorporate preference information into a setting with a baseline allocation, i.e, an *endowment*. For example, HIAS might start with the employment-maximizing outcome from Annie™ MOORE and then give refugees an option to submit preferences in order to improve their matches. A matching is then *individually rational* if every family is matched to a locality it weakly prefers to its endowment. In this case, because families have different sizes, a Pareto-efficient and individually rational matching cannot be achieved by only using the trading cycles that arise in the KTTC algorithm: It might be necessary to swap sets of families in order to guarantee feasible Pareto improvements. Indeed, it turns out that there is no strategy-proof mechanism that guarantees to find even a single Pareto improvement upon an endowment that is not Pareto-efficient (Proposition 2). We therefore relax Pareto efficiency by considering *Pareto-improving chains* in which swaps occur at the level of families. We define a matching to be *chain-efficient* if it cannot be improved by carrying out any Pareto-improving chain. We show that there does not exist any chain-efficient and strategy-proof mechanism (Theorem 1). In fact, a strategy-proof mechanism that Pareto improves upon an endowment that is not chain-efficient is not guaranteed to exist when there is more than one dimension (Theorem 2). In order to Pareto improve upon an endowment whenever possible, we introduce an algorithm, called Knapsack Top Trading Cycles with Endowment (KTTCE), which generalizes the KTTC algorithm. (Without an endowment, both algorithms are equivalent (Proposition 3).) The KTTCE mechanism is strategy-proof and can potentially Pareto improve upon the endowment by carrying out Pareto-improving chains (Theorem 3). If there is only one dimension and larger families have a higher priority, then the KTTCE algorithm is guaranteed to Pareto improve upon the endowment in a strategy-proof way (Theorem 4).

When priorities of localities also need to be taken into account, new trade-offs arise. In particular, stable matchings may not exist and determining whether a stable matching exists (or finding a stable matching when one exists) is a computationally intractable problem (McDermid and Manlove, 2010; Biró and McDermid, 2014).⁴ In our model, stable outcomes only exist under fairly strong conditions (e.g., if families are prioritized by sizes; see Proposition 4). To address the non-existence and computational shortcomings of stable matchings, we introduce an alternative solution concept called *weak envy-freeness*, which is based on envy-freeness (Sotomayor, 1996; Wu and Roth, 2018; Kamada and Kojima, 2020), but is less demanding. Envy-freeness eliminates envy, that is, it rules out matchings in which a family f is matched to locality ℓ but prefers locality ℓ' which hosts a family f' with a lower-priority than f at ℓ' . Weak envy-freeness is less stringent than envy-freeness. Weak envy-freeness allows family f to envy family f' as long as f' does not interfere with any higher-priority family at locality ℓ' . We show that a family-optimal weakly envy-free matching exists and can be found via a modification of the classical Deferred Acceptance (DA) algorithm (Gale and Shapley, 1962), which we call the Knapsack Deferred Acceptance (KDA) algorithm (Theorem 5). However, unlike in contexts such as school choice, this modification of the DA algorithm is manipulable because localities' choice functions (induced by the priorities and constraints) do not satisfy the cardinal monotonicity condition (Alkan, 2002; Hatfield and Milgrom, 2005). In fact, there is no family-optimal weakly envy-free and strategy-proof mechanism (Proposition 7), implying a trade-off between incentives and efficiency when the designer requires weak envy-freeness.⁵ We therefore develop a strategy-proof and weakly envy-free mechanism, called the Threshold Knapsack Deferred Acceptance (TKDA) algorithm (Theorem 6). We use refugee resettlement data from HIAS to test our mechanisms. We assume that localities rank refugees according to their likelihood of employment estimated by Trapp et al. (2020). As refugee preferences are currently not collected, we simulate different preference distributions. We find that (particularly when refugees' preferences are uncorrelated): (i) the KTTCE mechanism makes many families better off compared to the employment-maximizing endowment; (ii) the KDA algorithm is substantially more efficient than the TKDA algorithm; (iii) there are considerable efficiency gains from using weak envy-freeness criterion compared to envy-freeness in the KDA and TKDA algorithms in the presence of multidimensional constraints.

⁴We use the term *stable* in the sense of *pairwise stability* following the two-sided matching literature (Gale and Shapley, 1962). In matching markets with priorities, a stable matching is sometimes said to *eliminate justified envy* (Abdulkadiroğlu and Sönmez, 2003) or to be non-wasteful and *fair* (Balinski and Sönmez, 1999).

⁵We leave the study of mechanisms that satisfy weaker conditions on truth-telling incentives, such as Bayesian incentive compatibility (Ehlers, 2008), regret-freeness (Fernandez, 2017), or partial strategy-proofness (Mennle and Seuken, 2014), for future work.

Impact on resettlement practices To the best of our knowledge, no resettlement agency currently uses refugees’ preferences directly in determining refugee–locality assignments. However, as we elaborate in Section 2.1, this does not mean that resettlement agencies do not see value in taking preferences into account—and in fact, they have indicated an interest in doing so. But the constraints that resettlement agencies face mean that existing matching frameworks are not directly applicable to refugee resettlement. This paper provides the necessary theory to make preference-based refugee matching a reality. We study a matching model that accounts for (possibly multidimensional) knapsack constraints, highlight the trade-offs that will arise, and propose a number of mechanisms that sketch out parts of the allocative frontier.

Our KTTCE mechanism constitute a natural first step as it allows directly incorporating preferences into current practices. For example, if HIAS were to collect preferences, it could use the KTTCE mechanism in conjunction with the Annie™ MOORE system it currently uses: Annie™ MOORE would calculate an endowment and the KTTCE mechanism would identify and carry out mutually beneficial trades. Thus, each family would be matched, to a locality at least as preferred as the one they would receive under the current allocation scheme.

In the long run, an agency may wish to involve localities more closely in the matching process. Our mechanisms based on weak envy-freeness provide solutions that account for the priorities of localities alongside the preferences of refugees. Priority orders could come from a combination of employment probabilities (as calculated by Annie™ MOORE), administrative policies, and local residents’ preferences; accounting for these explicitly may help create goodwill from localities—who may in turn increase their capacity—or at least take advantage of the localities’ information about what constitutes a good match.

As for any matching market, which solution works best is likely to depend on specific circumstances. Only practice can tell us exactly what those might be in the refugee resettlement context. In that respect, our work offers a practically relevant contribution: it provides resettlement agencies with a motivating reason to collect preferences and evaluate how they would affect allocation under different matching rules.

Relationship to prior work. Matching markets for refugee resettlement were first proposed by [Moraga and Rapoport \(2014\)](#) as a part of a system of international refugee quota trading ([Schuck, 1997](#)). In the international context of matching refugees to countries, however, the refugee matching market is “thick”—any country can be expected to host any family up to its capacity—and can be reasonably modeled as a standard school choice problem ([Abdulkadiroğlu and Sönmez, 2003](#); [Jones and Teytelboym, 2017a](#)). [Jones and Teytelboym](#)

(2017b) informally introduced the idea of refugee resettlement matching in the national context and pointed out the knapsack constraints and the thinness of matching markets that arise on the local level. [Andersson and Ehlers \(2020\)](#) examine a market for allocating private housing to refugees in which landlords have preferences over the sizes of refugee families and over the native languages refugees speak. Our work is the first to offer a formal theory of preference-based matching for refugee resettlement.

Beyond refugee resettlement, our work draws upon and contributes to the applied literature on the design and implementation of complex matching mechanisms. The most famous example is the National Resident Matching Program (NRMP), in which residents may apply to jobs as couples ([Roth and Peranson, 1999](#); [Klaus and Klijn, 2005](#); [Klaus et al., 2007](#); [Haake and Klaus, 2009](#)). In this market, as in ours, stable outcomes may not exist. There are a number of algorithms that can find stable matchings in the couples model whenever they exist ([Echenique and Yenmez, 2007](#); [Kojima, 2015](#)) or find approximate solutions ([Nguyen and Vohra, 2018](#)). However, the structure of our problem is different to the matching with couples problem as the barriers to stability in our context arise from the constraints on the locality (hospital) side, rather than from the family (doctor) side (as in the couples problem). While stable matchings are computationally hard to find in our model (even when they exist), there are various weaker stability concepts that are computationally tractable ([Aziz et al., 2018](#)). Stable outcomes also do not exist in general in the market for trainee teachers in Slovakia and Czechia, where teachers are expected to teach two out of three subjects and schools have capacities for each subject ([Cechlárová et al., 2015](#)). Another difficult case for market design has been matching with minimum quotas, in which stable outcomes also typically do not exist ([Goto et al., 2014](#); [Fragiadakis et al., 2016](#)). In similar spirit to this paper, [Kamada and Kojima \(2020\)](#) consider many-to-one matching markets under general constraints. While their model is more general than ours, their results are independent of ours, and they focus on the structure of constraints that allow for an existence of a feasible, individually rational, and envy-free matchings. [Milgrom and Segal \(2020\)](#) study a dynamic auction with knapsack constraints, but they focus mainly on the properties of the classic deferred acceptance algorithm. Finally, [Nguyen et al. \(2019\)](#) consider a version of our model in which localities have cardinal preferences that arise from an integer optimization problem. However, [Nguyen et al. \(2019\)](#) focus on group-stable and near-feasible matchings and do not consider strategic issues which are key for preference-based refugee resettlement.

In some matching market design settings, such as school choice in New Orleans or housing allocation, stability is considered secondary to efficiency. In such cases, the Top Trading Cycles algorithm ([Balinski and Sönmez, 1999](#); [Abdulkadiroğlu and Sönmez, 2003](#)) or its modifications ([Pápai, 2000](#); [Dur and Ünver, 2015](#)) are used instead of stable mechanisms.

Pycia and Ünver (2017) show that in settings where agents have single-unit demands over objects, all Pareto-efficient mechanisms that cannot be manipulated by a group of agents can be represented in terms of a general class of Trading Cycles mechanisms. Pápai (2003, 2007) analyzes the difficulties of exchange with endowments and multiple goods. In our setting, families are only endowed with at most one locality; however, efficient and strategy-proof mechanisms are similarly hard to find. Abdulkadiroğlu et al. (2009) propose to take advantage of priority ties in school choice to improve upon the student-optimal stable matching. They show that any such improvement comes at the cost of losing strategy-proofness. Their negative result arises because the matching which they want to improve upon depends on reported preferences; on the other hand, our negative results rely solely on knapsack constraints as in our setting the endowment is independent of preferences.

Organization of the paper. The remainder of the paper is organized as follows. In Section 2, we describe the institutional context of refugee resettlement in the U.S.. We state our formal model in Section 3. In Section 4, we explain how two variations on the Top Trading Cycles algorithm can fully incorporate preferences of refugees. In Section 5, we propose solutions for the case where refugee preferences need to be balanced with priorities of the localities. In Section 6, we present simulations based on refugee resettlement data from HIAS. Section 7 is a conclusion. All proofs are in the Appendix. The Online Appendix provides additional results on the Threshold Knapsack Deferred Acceptance Algorithm (Online Appendix A), examples of how different algorithms work (Online Appendix B), further simulation results (Online Appendix C) and relationships between our model and previous models (Online Appendix D).

2 Institutional context

The Vietnam War and the evacuation of over 130,000 Vietnamese, Cambodian, and Laotian refugees to the U.S. in 1975 precipitated the Refugee Act of 1980, which created the federal Office for Refugee Resettlement. This Act standardized the resettlement process, set flexible annual quotas, and fixed funding for resettlement. Since then around three million refugees have been resettled to the U.S., mainly from southeast Asia and the former Soviet Union. Annual resettlement numbers fluctuate considerably, partly because they can be altered by executive action in response to crises and political will. In the past decade, around 70,000 refugees arrived annually to the U.S.; this rose to almost 78,761 in 2016. In September 2016, President Obama had committed to resettling at least 110,000 refugees in 2017; however, in January 2017 President Trump reduced the quota to 50,000 (although only 24,559 refugees

were eventually resettled to the U.S. in 2017 and 17,112 in 2018).

Refugees can apply for the U.S. resettlement program directly or be referred by the UNHCR (often while living in a refugee camp). The refugee resettlement program is managed by the U.S. Refugee Admissions Program (USRAP) which, alongside the UNHCR and the International Organization for Migration, identifies refugees, conducts security checks on all family members, and arranges travel;⁶ This process can take 18 to 24 months. Once a family has passed the security checks, it proceeds to medical checks and cultural orientation. Then the case is handed over to one of the nine U.S. Resettlement Agencies, which are responsible for matching the refugee family to a local community.⁷ Resettlement agencies resettle refugees all over the U.S., although the majority are initially placed in California, Florida, New York State, and Texas.

Refugees are allowed to list family members who live in the U.S., in which case they are almost certain to be reunited with them. But beyond that, resettlement agencies do not collect information about refugees' preferences over initial placements, and instead must make informed guesses about where refugees would fare well.

Resettlement agencies establish their own links to local communities, which we refer to as *localities* throughout the paper, that are willing to host refugees.⁸ Every year, agencies review the capacities of their localities to host refugees. Hosting commitments of localities run for the duration of the fiscal year (from October to October). Localities express a variety of hosting constraints to their agencies. The key constraint is the total number of refugees they are able or willing to host. The Department of State approves the quotas for every locality, which can be as high as 200 or as low as 20 refugees per year. Localities cannot exceed the annual quotas. There can also be additional constraints, for example on the number of refugees from certain countries or regions, the number of children or elderly members, or the number of refugees with a medical condition. Other (binary) constraints might include whether the community can support single-parent families or disabled refugees. Agencies assign refugees to localities roughly every fortnight. In order to balance the resettlement load, the fortnightly quota for each locality is set proportionally to that locality's annual

⁶USRAP consists of the Bureau of Population, Refugees and Migration (PRM) of the U.S. Department of State, the U.S. Citizenship and Immigration Services (USCIS) of the U.S. Department of Homeland Security, and the Office of Refugee Resettlement (ORR) of the U.S. Department of Health and Human Services (HHS)

⁷These resettlement agencies, also known as "voluntary agencies", are: Church World Service (CWS), Ethiopian Community Development Council (ECDC), Episcopal Migration Ministries (EMM), Hebrew Immigrant Aid Society (HIAS), International Rescue Committee (IRC), U.S. Committee for Refugees and Immigrants (USCRI), Lutheran Immigration and Refugee Services (LIRS), U.S. Conference of Catholic Bishops (USCCB), World Relief Corporation (WR).

⁸Resettlement agencies also coordinate the entire arrival process with the locality and ensure that housing and initial support facilities (e.g., airport pickup, pocket money, first meals) are ready when the refugee family arrives.

quota and is treated as a hard constraint (Trapp et al., 2020). Localities do not currently fix priorities over specific types of refugees beyond the constraints they express to agencies. Most localities commit to supporting refugees for their first year of resettlement, after which the refugees are more or less on their own.⁹

2.1 Introducing preferences and priorities into refugee resettlement

In 2018, HIAS became the first resettlement agency to adopt a matching system that attempts to maximize short-run refugee employment while meeting the constraints of localities. The system predicts employment likelihoods across refugee-locality matches using observable historical data and suggests matchings that aim to maximize the employment objective. While matching families and localities based on observable characteristics constitutes a significant improvement over random allocation or ad hoc procedures (Bansak et al., 2018; Trapp et al., 2020), there are still reasons to incorporate refugees’ preferences and localities’ priorities into the matching process.¹⁰

Refugees are likely to hold private information about their preferences that can affect the quality of matches. As Mark Hetfield, the CEO of HIAS has explained:

“Many Somali refugees initially settled around the country subsequently migrated to Lewiston, Maine. Lewiston has a weak economy but an established Somali community. Consequently, efforts to resettle these refugees elsewhere in the U.S. were less effective than they could have been. Their preferences should have been taken into account from the start” (see Roth (2015)).

Indeed, refugees’ private information might even be pertinent to maximizing observable outcomes, such as employment. As Hetfield points out, taking preferences into account may also help prevent internal migration—the movement of refugees away from their assigned localities soon after arrival—which localities want to avoid because they make substantial upfront investments in hosting refugees. On the other side of the match, resettlement agencies also miss information by not eliciting priorities from localities. For example, a locality cannot currently signal that it would prefer to resettle larger or smaller families, families from

⁹The federal government provides very limited support with a fixed grant of \$1,125 per refugee to cover the first 90 days of resettlement.

¹⁰Logistically, there is no obvious impediment to the collection of preferences and priorities. Resettlement agencies have time to elicit families’ preferences over localities after those families’ applications have been approved and before departure. Localities can be consulted about their priorities at the beginning of each fiscal year.

particular ethnic or linguistic groups, or families in which members have certain medical conditions.¹¹

Another reason for eliciting and responding to preferences is ethical. Refugees ought to be afforded the same agency as other citizens and be given some choice during one of the most consequential moments in their lives (Jones and Teytelboym, 2017b). It is arguably inconsistent to talk about what is “good” for refugees without seriously considering their preferences. Meanwhile, respecting priorities of localities can serve as an important fairness criterion for refugee families and give communities a sense of control over the resettlement process—which might corral goodwill and incentivize localities’ long-term participation.

Eliciting preferences of refugees over localities might appear to be a difficult task. Refugees face roughly the same number of options (16 at HIAS as of September 2020) as parents in a public school choice system—and these parents often face difficulties forming and expressing their preferences (Hastings et al., 2007; Corcoran et al., 2018). Moreover, refugees may not be able to determine precise preferences over individual localities because they lack necessary the information, such as an understanding of the localities’ labor markets and opportunities. To address this lack of information in resettlement matching, refugees could be asked to rank (or value) the *properties* of localities that are important to them: for example, proximity to a city, low crime, presence of a co-ethnic or a co-religious community (Jones and Teytelboym, 2017b). The resettlement agency can then combine refugees’ expressed preferences over characteristics with data on localities to infer a likely preference profile over localities (see, for example, Wiswall and Zafar (2017), who use the hypothetical choice methodology to estimate preferences for workplace attributes).

Eliciting localities’ priorities is more straightforward. Resettlement agencies might simply ask localities to list and weight the characteristics of families that they would prefer to resettle. Agencies might want to exclude some characteristics from the ranking process (e.g., race or religion), and expect localities to resettle any family that does not violate its capacity constraints. Alternatively, localities might want to rank families according to features that are difficult for localities themselves to estimate (e.g., likelihood of employment, see Bansak et al. (2018); Trapp et al. (2020)); in that case, the agency could estimate these characteristics on behalf of the localities and assign priority rankings accordingly.

¹¹A resettlement agency might put reasonable restriction on what criteria can be used for priority rankings. For example, a locality could be prevented from ranking refugees based on race or religion.

3 Model

There is a finite set of refugee *families* $f \in F$ and a finite set of *localities* $\ell \in L$. We assume that there exists a *null* locality $\emptyset \in L$ that represents being unmatched. There is a finite set D of *dimensions*. The *size* of a family f is the vector $(\nu_d^f)_{d \in D} \in \mathbb{Z}_{\geq 0}^{|D|} \setminus \{\mathbf{0}\}$.¹²

The vector $(\kappa_d^\ell)_{d \in D} \in \mathbb{Z}_{\geq 0}^{|D|}$ denotes the *capacity* of locality ℓ . We use the convention that $\kappa_d^\emptyset = +\infty$ for all dimensions d , i.e., that the null locality has infinite capacity. Locality ℓ can *accommodate* a set of families $G \subseteq F$ if $\sum_{g \in G} \nu_d^g \leq \kappa_d^\ell$ for all $d \in D$. Locality ℓ can *accommodate* a family f *alongside* $G \subseteq F \setminus \{f\}$ if ℓ can accommodate $G \cup \{f\}$. We assume that every family can be accommodated on its own by at least one locality other than the null.

A (feasible many-to-one) *matching* is a correspondence $\mu : F \cup L \rightrightarrows F \cup L$, such that for all $f \in F$, $\ell \in L$, and $d \in D$,

- (i) every family is matched to exactly one locality, i.e., $\mu(f) \in L$;
- (ii) every locality is matched to a subset of families, i.e., $\mu(\ell) \subseteq F$;
- (iii) a family is matched to a locality if and only if the locality is matched to the family, i.e., $\mu(f) = \ell$ if and only if $f \in \mu(\ell)$; and
- (iv) every locality can accommodate all the families matched to it, i.e., $\sum_{g \in \mu(\ell)} \nu_d^g \leq \kappa_d^\ell$.

The first three conditions are standard while condition (iv) ensures that no locality's capacity is exceeded. Figure 1 illustrates an example of a matching in a knapsack constraints. Our model generalizes a number of existing matching models with complex constraints (see Online Appendix D).

Families have *preferences* over localities. We denote by \succ_f the strict ordinal preference list of family f over L , and let $\succ = (\succ_f)_{f \in F}$ be the preference profile of families. We therefore write $\ell \succ_f \ell'$ to mean that f strictly prefers ℓ to ℓ' . We write $\ell \succeq_f \ell'$ to denote that either $\ell \succ_f \ell'$ or $\ell = \ell'$. We assume that every family's least preferred option is being unmatched.

Localities have exogenously fixed *priorities* over families. We let \triangleright_ℓ be the strict ordinal priority list of locality ℓ over families F and let $\triangleright = (\triangleright_\ell)_{\ell \in L}$ be the ordinal priority profile of the localities. We therefore write $f \triangleright_\ell f'$ to mean that f has a higher priority than f' at ℓ . We assume that localities prioritize families they can accommodate on their own: if ℓ can accommodate $\{f\}$ but not $\{f'\}$, then $f \triangleright_\ell f'$.

¹²We assume that every family takes up one unit of capacity in at least one dimension. This is without loss of generality as a family that does not take up any capacity can be matched to any locality without affecting any other family. Moreover, this is consistent with our application as, in practice, every family has at least one family member.

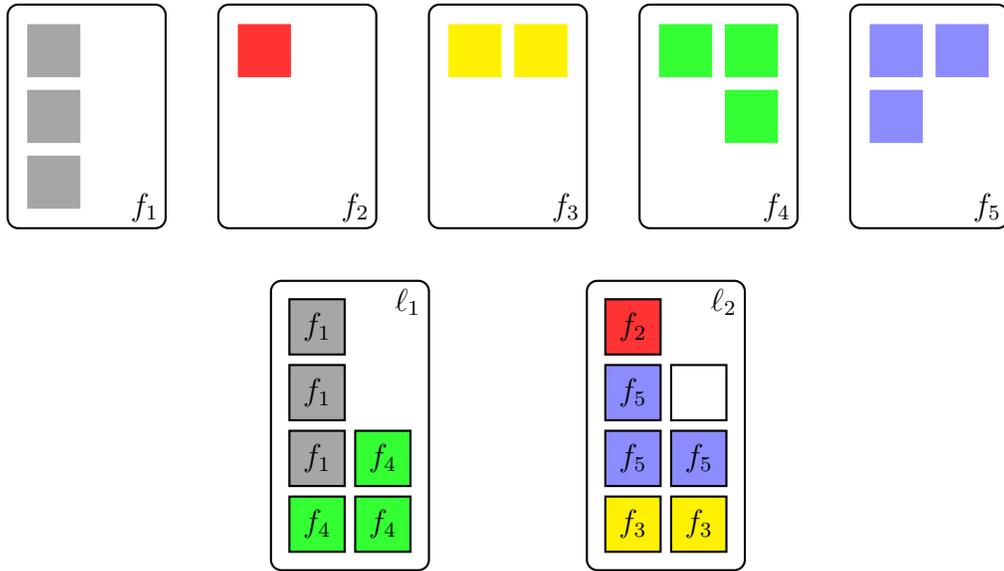


Figure 1: Matching in a market with two-dimensional constraints. There are five families f_1, \dots, f_5 , two localities l_1, l_2 , and two dimensions, represented by the left and right columns. The sizes of the families are $(3, 0)$, $(1, 0)$, $(1, 1)$, $(1, 2)$, and $(2, 1)$; for example, family f_1 's size is represented by 3 blocks in the first column and 0 blocks in the second column. The capacities of the localities are $(4, 2)$ and $(4, 3)$. In the matching pictured, families f_1 and f_4 are matched to locality l_1 and families f_2, f_3 , and f_5 are matched to locality l_2 . All of l_1 's capacity is used while l_2 has one unit capacity for the second dimension that remains unused.

A matching μ is *wasteful* if there exists a family $f \in F$ and a locality $\ell \in L$ such that $\ell \succ_f \mu(f)$ and ℓ can accommodate f alongside $\mu(\ell)$. We say that matching μ *Pareto dominates* matching μ' , denoted $\mu \succ \mu'$, if $\mu(f) \succeq_f \mu'(f)$ for all $f \in F$ and $\mu(f) \succ_f \mu'(f)$ for some $f \in F$. We write $\mu \succeq \mu'$ if μ *weakly Pareto dominates* μ' , that is if either $\mu \succ \mu'$ or $\mu = \mu'$. A matching μ is *Pareto-efficient* if there does not exist another matching μ' that Pareto dominates μ .

Fixing a set of families and their sizes, a set of localities and their capacities, and a profile of priorities, we define a (direct) *mechanism* as a function φ that takes as input a preference profile and outputs a matching. A mechanism φ is *strategy-proof* if for any $f \in F$ there does not exist a report of a preference list \succ'_f such that

$$\varphi(\succ'_f, \succ_{-f})(f) \succ_f \varphi(\succ)(f),$$

where $\varphi(\succ)(f)$ is the locality to which f is matched under the mechanism φ and the preference profile \succ . Strategy-proofness requires that refugee families cannot make themselves better off by misreporting their preferences over localities, irrespective of the reports of other families.

If a mechanism always selects a matching with a certain property, we refer to the mechanism as having that property. For example, if a mechanism always selects a non-wasteful matching, we call it a *non-wasteful mechanism*.

Throughout the paper, we describe all the mechanisms as algorithms: therefore, a mechanism takes a preference profile as an input and uses instructions from its corresponding algorithm to produce a matching. We say that a family is *permanently matched* (*permanently rejected*) to a locality at some step of the algorithm for mechanism φ if by that step it has been established that the family will (not) be matched to the locality in the matching outputted by φ .

3.1 Knapsack constraints in practice

In practice, there is a constraint on the total number of refugees that a locality can host. As we explained in Section 2, these quotas are approved by the Department of State and are treated as hard constraints. If we call this dimension d_1 and consider a family f , $\nu_{d_1}^f$ is the number of family members (e.g., if f is a dual-parent family with four children, then $\nu_{d_1}^f = 6$). For a locality ℓ , $\kappa_{d_1}^\ell$ is that locality's quota, i.e., the number of refugees that ℓ can host.

On top of this, resettlement agencies face additional constraints that are agreed on directly with the localities. In the case of HIAS, all of these constraints are currently binary.

For example, some localities are equipped to host single-parent families or refugees with particular disabilities or medical conditions while others are not. In addition, some localities only accept refugees speaking languages for which they have translators available. Additional dimensions are not required to model binary constraints: we can straightforwardly extend the model so that some family-locality pairs can simply be considered to be infeasible. Families would only have preferences over feasible localities and localities would only rank feasible families. All the results in the paper would go through with this minor modification.

HIAS currently uses only binary additional constraints so the special case of our model in which $|D| = 1$ is sufficient for introducing preferences into its current practice. While all our results and solution concepts are novel even in the one-dimensional case, there are reasons to think that our general model may be useful to apply preference-based matching more widely. First, until 2016, there was a second institutional constraint limiting the number of refugees from a given region. For example, a locality might be able to host up to fourteen refugees in total but no more than eight Syrians. Given how volatile the legal structures around refugee resettlement are, it is conceivable that additional institutional constraints may be added in the future and modeling such a situation would require adding a dimension for each of the constrained regions. Second, some of the binary constraints that HIAS uses do not necessarily need to be binary but are treated as such because managing multidimensional constraints manually is difficult. If the model allows for multidimensional constraints, it may be possible to relax some of the binary constraints on the matching. For example, a locality who has only one Arabic translator available may be able to take one Arabic-speaking family but not more, a situation for which binary constraints cannot cater but multidimensional constraints can. Third, the possibility of using multidimensional constraints could help localities express their ability to support different kinds of refugee families. For example, a constraint on the total number of children could reflect the capacity of schools and kindergartens and the constraint on the number of seniors could reflect the capacity of care homes (see Section 6.3).

From a practical perspective, it is important to distinguish constraints from preferences. Refugees have preferences over localities and we allow these to be arbitrary. Constraints are imposed by the localities (in agreement with resettlement agencies and the U.S. Department of State) which can specify a total number of refugees they are able to host as well as, possibly, additional constraints for specific categories of refugees or families. Therefore, independently of the families' preferences, any matching must satisfy all of these constraints.

As we mentioned in Section 2, while quotas on total number of refugees are set annually, refugees arrive over the course of the year so the matching takes place roughly every fortnight. The current practice for HIAS is to consider each static matching problem inde-

pendently: Every fortnight, there is a set of families to be matched to a set of localities and the total number of refugees that a locality can host is set proportionally to its remaining annual quota and treated as a hard constraint (Trapp et al., 2020).¹³ Adjusting capacity across matching periods over the course of the year or changing the frequency of the matchings (e.g., to monthly or quarterly) could in principle improve both the employment-based optimization (Trapp et al., 2020, Table 10) and preference-based matching. We focus on the static matching problem and leave dynamic capacity management for future work.

4 Targeting efficiency

In this section, we propose matching mechanisms that incorporate refugee preferences and give refugees an incentive to report their preferences over localities truthfully.

We show first that the Pareto-efficient and strategy-proof Top Trading Cycles (TTC) mechanism of Shapley and Scarf (1974) can be adapted to our setting in a natural way while preserving its properties. However, resettlement agencies already have existing processes for assigning refugees to localities—for example, HIAS matches refugees according to observable characteristics to maximize objectives such as the likelihood of employment (Bansak et al., 2018; Trapp et al., 2020). To make preference-based matching as easy as possible to integrate with existing systems, we extend our TTC-based approach to allow resettlement agencies to use their initial allocations—which we refer to as an *endowment*—as a baseline and allow refugees to express preferences in order to improve upon that baseline.

We consider mechanisms that Pareto improve upon the endowment in the sense that every refugee family is assigned to a locality that it weakly prefers to its endowment. Pareto improvements upon an endowment make it possible to use refugees’ preferences to identify welfare gains over and above the assignments that would be selected based on observables. Such Pareto improvements also allow, in principle, resettlement agencies to institute distributional goals or other minimal welfare guarantees through the choice of the endowment. We show that finding Pareto improvements upon an endowment presents many challenges in our environment; even so, we obtain a strategy-proof mechanism based on the TTC algorithm that can achieve Pareto improvements.

¹³Since the decisions about the timing of arrivals are made by the U.S. government, there is little or no scope for refugees to manipulate the timing of arrivals (as in Caspari, 2019). Moreover, since the annual quotas are fixed at the beginning of the fiscal year and then managed by the resettlement agencies, there is no short-run scope for quota manipulation by localities (as in Sönmez, 1997).

4.1 The Knapsack Top Trading Cycles mechanism

Our first mechanism, described in Algorithm 1, is an extension of the TTC mechanism to matching with (possibly multidimensional) knapsack constraints.

In the first round, each family f *points* at its most preferred locality ℓ that can accommodate f and each locality ℓ' *points* at the highest-priority family f' that ℓ' can accommodate. There must be at least one (directed) cycle in the graph in which nodes are labeled with families and localities and “pointing” is represented by directed edges. Note that each family and each locality is in at most one cycle. Every family in a cycle is permanently matched to the locality at which it is pointing.

In the subsequent rounds, every locality permanently rejects every family that cannot be accommodated alongside families that are already permanently matched to that locality. Every family then points at its most preferred locality from which the family has not yet been permanently rejected; every locality points at its highest-priority family that the locality has not yet permanently rejected. Once again, there is at least one cycle and in each cycle we permanently match every family to the locality at which that family is pointing. The algorithm terminates in a finite number of steps because at least one family is permanently matched in each step of the algorithm.

Proposition 1. *The KTTC mechanism is strategy-proof and Pareto-efficient.*

The Knapsack Top Trading Cycles (KTTC) algorithm generalizes the classical TTC algorithm used in school choice by taking into account knapsack constraints. The KTTC mechanism is Pareto-efficient because a family continues pointing at its most preferred locality until it can no longer be accommodated alongside families that are permanently matched to it. The key reason for strategy-proofness is that a family can only be permanently matched by being involved in a cycle. A family can only create such a cycle by pointing at any locality in a pointing sequence that terminates with the family; the pointing sequence continues to exist until the family is permanently matched. Therefore, every family has an incentive to continue pointing at its most preferred locality that has not permanently rejected it yet.

Note that Proposition 1 does not include any properties of the KTTC mechanism that pertain to the priorities of the localities. Indeed, the priorities are only used in the KTTC mechanism in order to determine the order in which localities point. However, strategy-proofness and efficiency of the KTTC mechanism do not depend on the pointing order of the localities.¹⁴

¹⁴If all localities have the same priority order, then the KTTC mechanism collapses to the serial dictatorship mechanism.

Algorithm 1: KNAPSACK TOP TRADING CYCLES (KTTC)

Initialize the *current matching* μ^1 such that $\mu^1(f) = \emptyset$ for all $f \in F$. No families are *permanently matched*.

Round $i \geq 1$

Starting with μ^i , every locality ℓ *permanently rejects* all families that ℓ cannot accommodate alongside $\mu^i(\ell)$.

Every family f that is not permanently matched points at its most preferred locality among those that have not permanently rejected f .

Every locality ℓ points at the highest-priority family that has not been permanently matched and that ℓ has not permanently rejected. (If no such family exists, then ℓ does not point.)

At least one cycle appears and every family and every locality is involved in at most one cycle. Update the current matching to μ^{i+1} by permanently matching every family in a cycle to the locality at which it is pointing (this could be \emptyset).

If all families are permanently matched, end and output μ^{i+1} . Otherwise continue to Round $i + 1$.

4.2 Improving efficiency from an endowment

Consider an exogenous matching μ^E , which we refer to as the *endowment*. We say that a matching μ is *individually rational* if $\mu \succeq \mu^E$, that is if μ weakly Pareto dominates the endowment.¹⁵ We will also refer to locality $\mu^E(f)$ as *family f 's endowment* and to families $\mu^E(\ell)$ as *locality ℓ 's endowment*. In the school choice setting, i.e., when $|D| = 1$ and $\nu_d^f = 1$ for all f , the TTC mechanism finds an individually rational and Pareto-efficient matching by carrying out one cycle at a time.¹⁶ In contrast, in a setting with knapsack constraints, carrying out one cycle at a time may not achieve Pareto-efficiency because some Pareto improvements may require families to swap in groups. For example, two “small” families in one locality might be able to swap simultaneously with a “large” family in another locality, but none of the pairwise swaps between a “small” family and the “large” family would be feasible. We therefore limit the set of Pareto improvements that can be executed.

Definition 1. Given a matching μ , a *Pareto-improving chain* is a sequence

$$(f_1, \ell_1, f_2, \ell_2, \dots, f_n, \ell_n)$$

of distinct families and localities such that:

- $\ell_1 \succ_{f_1} \mu(f_1)$;
- for all $i = 2, \dots, n$,
 - ℓ_{i-1} can accommodate f_{i-1} alongside $\mu(\ell_{i-1}) \setminus \{f_i\}$,
 - $\ell_i \succ_{f_i} \ell_{i-1} = \mu(f_i)$; and
- ℓ_n can accommodate f_n alongside $\mu(\ell_n) \setminus \{f_1\}$.

In any Pareto-improving chain, family f_1 moves to locality ℓ_1 which f_1 prefers to its current locality. Locality ℓ_1 , in turn, must be able to accommodate f_1 alongside all families in $\mu(\ell_1)$ except for f_2 which leaves locality ℓ_1 for a more preferred locality ℓ_2 . The Pareto-improving chain continues with f_3 moving from ℓ_2 to ℓ_3 and so on. The Pareto-improving chain terminates in one of two ways: either (i) the Pareto-improving chain is “open” and no family leaves the last locality, i.e., $\ell_n \neq \mu(f_1)$ or (ii) the Pareto-improving chain is “closed” and f_1 leaves the last locality, i.e., $\ell_n = \mu(f_1)$.

Definition 2. A matching is *chain-efficient* if it does not have any Pareto-improving chain.

¹⁵Our meaning from individual rationality is in line with the rest of the social choice literature here as the endowment, rather than not being matched, now plays the role of the outside option.

¹⁶The KTTTC mechanism collapses to the TTC mechanism in the school choice setting.

Chain efficiency constitutes a relaxation of Pareto efficiency because it only requires the elimination of Pareto-improving chains, which form a subset of all possible Pareto improvements. In school choice, Pareto efficiency is equivalent to chain efficiency and the Pareto-efficient TTC mechanism is strategy-proof. In our setting with knapsack constraints, however, there might be a matching that Pareto dominates a chain-efficient matching if there are groups of families that could participate in a Pareto-improving swap that is not a Pareto-improving chain. Moreover, even chain-efficient mechanisms are not strategy-proof.

Theorem 1. *There is no strategy-proof, individually rational, and chain-efficient mechanism.*

The proof of Theorem 1 requires only one dimension and the largest family size is two. The intuition is that different Pareto-improving chains can interfere with each other, thereby giving families an opportunity to select into the Pareto-improving chain they prefer by manipulating their preferences. As Pareto efficiency implies chain efficiency, Theorem 1 directly implies the following result.

Corollary 1. *There is no strategy-proof, individually rational, and Pareto-efficient mechanism.*

More generally, Theorem 1 and Corollary 1 imply a trade-off between efficiency and strategy-proofness when the designer wants to Pareto improve upon an endowment. This trade-off does not exist in school choice and constitutes a direct consequence of the fact that families have different sizes. This leaves open an important question: can a strategy-proof mechanism guarantee even a single Pareto improvement upon an endowment? To formalize this idea, we say that a mechanism φ *Pareto improves* upon an endowment μ^E , if $\varphi(\succ) \succ \mu^E$, that is if the mechanism returns a matching that Pareto dominates the endowment. This definition strengthens individual rationality by ruling out the case where the mechanism returns the endowment.

Proposition 2. *There is no strategy-proof mechanism that Pareto improves upon every endowment that is not Pareto-efficient.*

By definition, a mechanism cannot Pareto improve upon a Pareto-efficient endowment. Proposition 2 implies that, even if it were possible to Pareto improve upon the endowment, one might not be able to do so without giving families an incentive to misrepresent their preferences. The intuition behind the proof of Proposition 2 is similar to the proof of Theorem 1.

As chain efficiency is less stringent than Pareto efficiency, one may still hope to be able to Pareto-improve upon endowments that are not chain-efficient in a strategy-proof way—but even this turns out to be impossible in the presence of multidimensional knapsack constraints.

Theorem 2. *If $|D| > 1$, there is no strategy-proof mechanism that Pareto improves upon every endowment that is not chain-efficient.*

Theorem 2 considers *any* kind of Pareto improvements, whether or not they are Pareto-improving chains. Therefore, the result directly implies that, when $|D| > 1$, it may not be possible to find any Pareto-improving chains—even if they exist—without giving families an incentive to misrepresent their preferences. This means that strategy-proofness may preclude *all* trade in matching markets with multidimensional knapsack constraints.

While Theorem 1 and Proposition 2 hold even when there is only one dimension, the impossibility result in Theorem 2 relies on the failure of the size monotonicity condition, which can only occur when there is more than one dimension. In the remainder of this section, we show how we can adapt the KTTC mechanism in order to overcome the impossibility result in Theorem 2 whenever sizes are monotonic (e.g., when there is only one dimension).

4.3 KTTC with Endowment

We now present an extension of the KTTC mechanism which attempts to Pareto improve upon an endowment. As in the KTTC algorithm, the KTTC with Endowment (KTTCE) algorithm (Algorithm 2) looks for trading cycles and families point at localities which they prefer to their endowment. However, the KTTCE algorithm checks whether the cycles that appear are feasible.¹⁷ In general, endowments can cause trading cycles to be infeasible.¹⁸ If trading cycles are feasible, we match families to the localities they are pointing at in the cycles just as in the KTTC algorithm. The key step—the Rejection Stage—deals with the case when none of the cycles are feasible. In the Rejection Stage, we pick a family f at random or according to some exogenous rule among those at which at least one locality is pointing. Family f is permanently rejected by every locality ℓ where f cannot be accommodated alongside families that are currently matched to ℓ except for some family at which ℓ is pointing (if there is such a family). Since f is not necessarily rejected by all localities, the Rejection Stage leaves an opportunity for f to be involved in a feasible cycle in a subsequent round of the algorithm. It is worth noting that any feasible cycle found by the KTTCE algorithm corresponds to at least one Pareto-improving chain.¹⁹ Therefore, the KTTCE

¹⁷This is not particularly challenging from a computational point of view as it simply requires to verify whether each locality in the cycle can replace the family pointing at the locality by the family at which the locality points. Thus, the KTTCE algorithm works in polynomial time.

¹⁸If $|D| = 1$ and $v_d^f = 1$ for all $f \in F$ and all $d \in D$ all trading cycles are feasible. Moreover, without endowments all trading cycles are also feasible.

¹⁹In particular, a feasible cycle can be broken into multiple “open” Pareto-improving chains if some localities point at families not in their endowment.

Algorithm 2: KTTC WITH ENDOWMENT (KTTCE)

Initialize the *current matching* μ^1 such that $\mu^1(f) = \mu^E(f)$ for all $f \in F$. No families are *permanently matched*.

Round $i \geq 1$

Every locality ℓ *permanently rejects* all families that ℓ cannot accommodate alongside families that are permanently matched to ℓ .

Every family f that is not permanently matched points at its most preferred locality among those that have not permanently rejected f .

Every locality ℓ points at the highest-priority family that has not been permanently matched and that ℓ has not permanently rejected. (If no such family exists, then ℓ does not point.)

At least one cycle appears and every family and every locality is involved in at most one cycle. Label the families and localities in any such cycle $f_1 \rightarrow \ell_1 \rightarrow f_2 \rightarrow \ell_2, \dots, f_n \rightarrow \ell_n \rightarrow f_1$.

A cycle is *feasible* if, for all $j = 1, \dots, n$, ℓ_j can accommodate f_j alongside $\mu^i(\ell_j) \setminus \{f_{j+1}\}$ (letting $f_{n+1} = f_1$).

If one or more cycles are feasible, continue to the *Matching Stage*. Otherwise, continue to the *Rejection Stage*.

Matching Stage: Update the current matching to μ^{i+1} by matching every family in a feasible cycle to the locality at which it is pointing (this could be \emptyset); all these families become permanently matched.

If all families are permanently matched, end and output μ^{i+1} . Otherwise continue to Round $i + 1$.

Rejection Stage: Pick one family f (at random or according to some exogenous rule) at whom at least one locality is pointing. Permanently reject f from all localities to which f cannot be matched, i.e., ℓ permanently rejects f if ℓ cannot accommodate f alongside $\mu^i(\ell) \setminus \{f'\}$ (where f' is the family at which ℓ is pointing).

If f is permanently rejected by either the locality at which f is pointing or by a locality pointing at f , let $\mu^{i+1} = \mu^i$ and continue to Round $i + 1$. Otherwise, pick another family that has not been picked yet and repeat the Rejection Stage.

algorithm attempts to improve upon the endowment by carrying out successive Pareto-improving chains (each of which may be “open” or “closed”).

Theorem 3. *The KTTCE mechanism is strategy-proof and individually rational.*

The KTTCE mechanism preserves strategy-proofness because (i) in the Matching Stage, families point at their most preferred localities (as in KTTC), and (ii) in the Rejection Stage, the permanent rejections from localities do not depend on reported preferences.

Effectively, the KTTCE algorithm adds the Rejection Stage to each round of the KTTC algorithm in order to deal with infeasible cycles created by the endowment. The following proposition formalizes this point—when all families are endowed with the null locality, the outcomes of the KTTC and KTTCE algorithms coincide.

Proposition 3. *If $\mu^E(f) = \emptyset$ for all $f \in F$, then $\mu^{KTTC} = \mu^{KTTCE}$.*

We illustrate the KTTCE algorithm with an example in Online Appendix B.2.

As Theorem 2 shows, whenever $|D| > 1$, no strategy-proof mechanism (e.g., KTTCE) can be guaranteed to find any Pareto improvements upon an endowment that is not chain-efficient. Therefore, in general, the KTTCE mechanism might simply output the endowment even if Pareto-improving chains exist.

4.4 Guaranteeing Pareto improvements in the KTTCE mechanism

We now introduce a condition on family sizes, which guarantees that there exist priority profiles for which the KTTCE mechanism improve upon every endowment that is not chain-efficient.

Definition 3. Sizes are *monotonic* if for any two families f and f' and any two dimensions d and d' , $\nu_d^f > \nu_d^{f'}$ implies $\nu_{d'}^f \geq \nu_{d'}^{f'}$.

When $|D| = 1$, sizes are always monotonic; however, when $|D| > 1$, the size monotonicity condition states that whenever f is larger than f' in one dimension, then f is weakly larger in all other dimensions.

We next provide a class of priority profiles which guarantees that the KTTCE mechanism improves upon any endowment that is not chain-efficient.

Definition 4. Let sizes be monotonic. A priority profile is *lexicographic* if, for every $f, g \in F$ and every $\ell \in L \setminus \{\emptyset\}$,

- families in ℓ 's endowment have a higher priority, i.e., $f \in \mu^E(\ell)$ and $g \notin \mu^E(\ell)$ implies $f \triangleright_\ell g$; and
- within ℓ 's endowment, larger families have a higher priority, i.e., $f, g \in \mu^E(\ell)$ and $\nu_d^f > \nu_d^g$ for some $d \in D$ imply $f \triangleright_\ell g$.

Lexicographic priorities imply that each locality prioritizes all families in its endowment over those that are not and that, among the families in its endowment, the locality prioritizes families in decreasing order of sizes. There are no restrictions about how each locality ranks any two families with the same size or any two families not in its endowment. Note that lexicographic priorities are only well-defined when sizes are monotonic. If sizes are not monotonic, then there exist two families, say f and g , such that $\nu_{d_1}^f > \nu_{d_1}^g$ and $\nu_{d_2}^f < \nu_{d_2}^g$. Lexicographic priorities dictate that every locality prioritize f over g and prioritize g over f , a contradiction.

If sizes are monotonic and priorities are lexicographic, at the start of the KTTCE algorithm every locality ℓ points at the largest family in its endowment. Suppose that ℓ permanently rejects a family f that is not in its endowment. This means that f cannot be accommodated at ℓ even when the *largest* family in ℓ 's endowment has been removed. Therefore, in the endowment, there does not exist any Pareto-improving chain in which f moves to ℓ . If the outcome of the KTTCE algorithm is the endowment, then f has been permanently rejected by every preferred locality. This means that the endowment had no Pareto-improving chains in which f moves to any of the preferred localities and f is not involved in any Pareto-improving chain in the endowment. By extending this argument to every family, we can see that the KTTCE algorithm only returns the endowment whenever the endowment is chain-efficient.

Theorem 4. *If sizes are monotonic and priorities are lexicographic, then the KTTCE mechanism Pareto improves upon every endowment that is not chain-efficient.*

The KTTCE mechanism takes priorities as an input to determine the family to which each locality points (if any) in every round. However, properties presented in this section do not depend on priorities. Therefore, if sizes are monotonic but priorities are not lexicographic, it is possible to modify the pointing order in the KTTCE algorithm by constructing the pointing order from lexicographic priorities. This does not affect the mechanism's properties—individual rationality and strategy-proofness—but ensures that the KTTCE mechanism finds a Pareto-improving chain as long as one exists. This observation implies the following corollary to Theorem 4.

Corollary 2. *If sizes are monotonic, then there exists a strategy-proof mechanism that Pareto improves upon every endowment that is not chain-efficient.*

Recall that whenever there is only one dimension, sizes are always monotonic ensuring the existence of a strategy-proof mechanism—i.e., KTTCE with an adjusted pointing order—that Pareto improves upon any endowment. Corollary 2 therefore contrasts the impossibility result in Theorem 2 by highlighting the possibility of strategy-proof Pareto-improvement upon endowments that are not chain-efficient in the case when there is only one dimension. Moreover, Corollary 2 is relevant for current practice: HIAS currently uses Annie™ MOORE with only one dimension—the only constraint is on the total number of refugees—in its resettlement processes (Trapp et al., 2020). In Section 6, we simulate realistic resettlement settings with one and three dimensions and show that the KTTCE mechanism finds Pareto improvements even when Theorem 4 does not apply. We leave open the question of finding the most efficient, strategy-proof, and individually rational mechanism.

5 Accounting for priorities

As we have argued in this paper, preferences of refugees are central to designing matching mechanisms for refugee resettlement. However, there are good reasons for taking priorities of localities seriously as well. Mechanisms in Section 4 do not guarantee to satisfy the priorities of localities. In this section, we offer mechanisms that respect priorities of localities in addition to the preferences of refugee families.

5.1 (Non-)existence of stable matchings

Denote by $\widehat{F}_\ell^f = \{g \in F : g \triangleright_\ell f\}$ the set of families with a higher priority than family f at locality ℓ . A common solution concept for balancing preferences and priorities is (pairwise) *stability* (Roth, 1984a; Abdulkadiroğlu and Sönmez, 2003).

Definition 5. A matching μ is *stable* if there is no $f \in F$ and $\ell \in L$ such that

- (i) f prefers ℓ to its current match, i.e., $\ell \succ_f \mu(f)$; and
- (ii) ℓ can accommodate f alongside $\widehat{F}_\ell^f \cap \mu(\ell)$, i.e., all families matched to ℓ with a higher priority than f .

In words, family f and locality ℓ are a *blocking pair* in a matching μ if f prefers ℓ to its current match and it is possible to accommodate f in ℓ without removing any higher-priority family. A matching is stable if it does not have any blocking pairs. Our definition extends the

concept of stability to a setting with possibly multidimensional knapsack constraints. Our definition is in line with the way stability is defined in similar models (see, e.g., [McDermid and Manlove \(2010\)](#); [Biró and McDermid \(2014\)](#); [Delacrétaz \(2019\)](#)). If $|D| = 1$ and $\nu_{d_1}^f = 1$ for all $f \in F$, [Definition 5](#) collapses to the “elimination of justified envy” used in school choice and other object allocation settings ([Abdulkadiroğlu and Sönmez, 2003](#)).

While stable matchings always exist in school choice models, they do not exist in ours (even when $|D| = 1$).²⁰ In fact, determining whether a stable matching exists in our model (even when $|D| = 1$) is a computationally intractable problem ([McDermid and Manlove, 2010](#)): The running time of an algorithm that can be guaranteed to find a stable matching or proves that none exists increases exponentially with the problem size.²¹ This can be an impediment to practical applications in large matching markets.²²

However, stable outcomes are guaranteed to exist in a special case of our model where sizes are monotonic and the priorities of localities are sufficiently similar.

Definition 6. A priority profile \triangleright is *aligned* if for any $f, g \in F$ such that $\nu^f \neq \nu^g$ and any $\ell, \ell' \in L \setminus \{\emptyset\}$, $f \triangleright_{\ell} g$ if and only if $f \triangleright_{\ell'} g$.

The aligned priorities condition generalizes the second part in the definition of lexicographic priorities ([Definition 4](#)). Under the aligned priorities condition, any two families with different sizes are ranked identically by all localities, but families with the same size can be ranked arbitrarily. The case of identical priorities is therefore a special case of the aligned priorities condition. If sizes are monotonic, aligned priorities also include the case where all localities give a higher priority to larger families and, symmetrically, the case where all localities give a higher priority to smaller families.

Proposition 4. *If sizes are monotonic and the priority profile is aligned, then a stable matching exists.*

The monotonicity of sizes and alignment of the priority profile ensures that the set of families can be partitioned into $\{F_1, F_2, \dots, F_n\}$ such that for any $i = 1, \dots, n$, all families in F_i have the same size and for all $j < i$, all families in F_j have a higher priority for all localities than all families in F_i . A stable matching then can be obtained in polynomial time by running sequentially the (family-proposing) Deferred Acceptance algorithm for each

²⁰In [Online Appendix B.1](#), we present an example of a market in which a stable matching does not exist.

²¹The decision problem of the existence of stable matchings in our setting is NP-complete, meaning there is no known efficient (e.g., with polynomial running time in the number of families or localities) method of solving it. Finding a stable matching, like solving the knapsack problem, is therefore NP-hard, i.e., as hard as the hardest computational problems. Verifying whether a particular matching is stable only involves checking all possible blocking pairs which is a simple computational problem.

²²[Delacrétaz et al. \(2016\)](#) present an algorithm to find a stable matching whenever one exists.

subset, starting with F_1 . In a school choice setting, all families have the same size ($\nu^f = 1$ for all $f \in F$) and priority profile is (trivially) aligned; therefore the existence of stable matchings follows immediately from Proposition 4. As in school choice, stable matchings under size monotonicity and priority alignment can be found in polynomial time in our model.

5.2 Envy-free matchings

In the refugee resettlement context, priority alignment might be a strong assumption. In general, we should expect localities' priorities to be heterogeneous. For example, even if all localities use the same objective function, such as employment, refugee families' likelihoods of employment might vary substantially across localities. The possible non-existence of stable matchings and the computational challenges involved in finding them motivates us to consider alternative solution concepts. The key issue we will face is how to trade off respect for priorities against tolerating wasted capacity. Indeed, as Delacrétaz (2019) shows: if waste must be eliminated, then there may exist a blocking pair (f, ℓ) where an arbitrarily large number of units of ℓ are assigned to families with a lower priority than f . In the context of refugee resettlement, locality goodwill might be important; therefore respecting priorities might often take center stage. Furthermore, the underuse of capacity may well be tolerable because it can be used for the next cohort of resettled refugee families. In the refugee resettlement context, therefore, it seems reasonable to consider matching mechanisms that tilt the balance somewhat in favor of respecting priorities rather than eliminating waste.

We now introduce *envy-free* matchings, which respect the priorities of localities, but introduce possible underuse of the localities' capacity.

Definition 7. Given a matching μ , $f \in F$ *envies* family $f' \neq f$ if

- (i) f prefers f' 's locality to its current match, i.e., $\mu(f') \succ_f \mu(f)$; and
- (ii) f has a higher priority at f' 's locality, i.e., $f \triangleright_{\mu(f')} f'$.

Definition 8. A matching μ is *envy-free* if, under μ , no family envies another family.

Envy-freeness ensures that priorities are fully respected (as in a stable matching), but it tolerates waste.²³ In our setting, waste can occur when a “small” family f prefers locality ℓ to its current match and could be accommodated by ℓ alongside $\mu(\ell)$, but a “larger” family with a higher priority which also prefers ℓ to its current match cannot be accommodated at

²³The concept of envy-freeness has been studied in two-sided matching theory (Sotomayor, 1996; Wu and Roth, 2018; Kamada and Kojima, 2020). While Wu and Roth (2018) use the term “envy-free”, Sotomayor (1996) and Kamada and Kojima (2020) respectively call the concept “simple” and “fair”.

ℓ . If the “smaller” family were matched to ℓ , it would be envied by the “large” family (even though the “large” family could not be accommodated at ℓ).

The existence of an envy-free matching is immediate—an empty matching in which all families are matched to the null is envy-free. Consequently, it is straightforward to show the existence of an envy-free matching that is not Pareto-dominated by another envy-free matching. It is, however, not obvious whether an optimal—from the point of view of families—envy-free matching nevertheless exists. Formally, an envy-free matching μ is *family-optimal* if $\mu \succeq \mu'$ for every envy-free matching μ' , i.e., μ weakly Pareto dominates every other envy-free matching.

Proposition 5. *There exists a unique family-optimal envy-free matching.*

The existence of the family-optimal envy-free matching provides one natural solution to balancing refugees’ preferences and localities’ priorities.²⁴ As we argued above; however, the family-optimal envy-free matching may lead to waste.

We now show that envy-freeness is too restrictive and unnecessarily wasteful when there are multiple dimensions and that certain kinds of innocuous priority violations should be allowed in order to improve refugee family welfare. We illustrate this point with an example.

Example 1 (Inadequacy of envy-freeness in matching with multidimensional constraints). There are three families f_1 , f_2 , and f_3 and one locality ℓ_1 . The priority list of ℓ_1 is $f_1 \triangleright_{\ell_1} f_2 \triangleright_{\ell_1} f_3$. There are two dimensions and the sizes and capacities are displayed below:

$$\nu = \begin{matrix} & f_1 & f_2 & f_3 \\ \begin{matrix} d_1 \\ d_2 \end{matrix} & \begin{pmatrix} 1 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \end{matrix} \quad \kappa = \begin{matrix} & \ell_1 \\ \begin{matrix} d_1 \\ d_2 \end{matrix} & \begin{pmatrix} 1 \\ 1 \end{pmatrix} \end{matrix}.$$

The family-optimal envy-free matching assigns f_1 to ℓ_1 and leaves the other two families unmatched. Envy-freeness precludes assigning f_3 to ℓ_1 because f_2 would then envy f_3 . However, f_2 and f_3 are not really in competition since the units they require are in different dimensions. Therefore, a matching in which both f_1 and f_3 are matched to ℓ_1 is more efficient and appears to have no meaningful priority violations. \square

5.3 Weakly envy-free matchings

Example 1 shows that even the family-optimal envy-free matching can introduce unnecessary waste by requiring removal of innocuous priority violations. We therefore introduce a solution

²⁴In recent work, [Kamada and Kojima \(2020\)](#) extend Proposition 5 by showing the existence of a family-optimal envy-free matching in a more general setting.

concept that relaxes envy-freeness by allowing families to be matched to a locality as long as they do not compete with higher-priority families.

Recall that a locality ℓ can accommodate a family f alongside set of families $G \subseteq F \setminus \{f\}$ if $\nu_d^f + \sum_{g \in G} \nu_d^g \leq \kappa_d^\ell$ for all $d \in D$. We begin by relaxing this definition.

Definition 9. Locality $\ell \in L$ can *weakly accommodate* family $f \in F$ alongside $G \subseteq F \setminus \{f\}$ if, for all $d \in D$,

$$\text{either } \nu_d^f = 0 \quad \text{or} \quad \nu_d^f + \sum_{g \in G} \nu_d^g \leq \kappa_d^\ell.$$

Definition 9 relaxes the original concept of accommodation by only taking into account the dimensions in which f will take at least one unit of capacity. If the aggregate sizes of families in G exceed the capacity of locality ℓ for some dimension, f cannot be accommodated alongside G ; however, it may still be possible to weakly accommodate f if f 's size in that dimension is zero. If $|D| = 1$, weak accommodation is equivalent to accommodation; however, if $|D| > 1$, ℓ may be able to weakly accommodate f alongside G but unable to accommodate f alongside G . In fact, ℓ may even be able to weakly accommodate f alongside G even if ℓ cannot accommodate the families in G . In Example 1, ℓ_1 cannot accommodate f_3 alongside f_1 and f_2 because it cannot even accommodate $\{f_1, f_2\}$. However, ℓ_1 can weakly accommodate f_3 alongside f_1 and f_2 because $\nu_{d_1}^{f_3} = 0$.

We are now in a position to define this section's main solution concept.

Definition 10. Given a matching μ , family $f \in F$ *strongly envies* $f' \neq f$ (with $\ell' = \mu(f')$) if

- (i) f prefers f' 's locality to its current match, i.e., $\ell' \succ_f \mu(f)$;
- (ii) f has a higher priority at ℓ' , i.e., $f \triangleright_{\ell'} f'$; and
- (iii) ℓ' cannot weakly accommodate f' alongside all families with a higher priority than f' at ℓ' that weakly prefer ℓ' to their current matches, i.e., $\{g \in F : g \triangleright_{\ell'} f' \text{ and } \ell' \succeq_g \mu(g)\}$.²⁵

Definition 11. A matching μ is *weakly envy-free* if no family strongly envies another family.

Weak envy-freeness relaxes envy-freeness by allowing some innocuous priority violations. More precisely, under weak envy-freeness, family f may envy family f' so long as f' 's size in all dimensions that prevent f from being matched to $\mu(f')$ is zero. All envy-free matchings

²⁵In fact, if there exist a family f' and a locality ℓ' satisfying condition (iii), then there exists a family $f \in \{g \in F : g \triangleright_{\ell'} f' \text{ and } \ell' \succeq_g \mu(g)\}$ for which conditions (i) and (ii) hold. We include conditions (i) and (ii) in Definition 10 in order to emphasize the relationship between envy-freeness and weak envy-freeness.

are weakly envy-free but the converse does not hold, even when $|D| = 1$.²⁶ In Example 1, the matching that assigns both f_1 and f_3 to ℓ_1 is weakly envy-free, but not envy-free. The reason for this is that even though f_2 has a higher priority than f_3 at ℓ_1 , f_3 's size in the first dimension is zero, i.e., $\nu_{d_1}^{f_3} = 0$. Therefore, f_2 envies f_3 , but does not strongly envy f_3 , when f_3 is matched to ℓ_1 .

Let us illustrate how the particular way in which weak envy-freeness relaxes envy-freeness might be relevant in refugee resettlement.

Example 2. There is one locality ℓ (other than the null) that can accommodate up to 99 children overall and one refugee with a rare medical condition: $\kappa^\ell = (99, 1)$.²⁷ There are two families f_1 and f_2 that do not have any children but each have one member with the rare medical condition— $\nu^{f_1} = \nu^{f_2} = (0, 1)$ —and 100 families $f_3 - f_{102}$ that each have one child but no individual with the rare medical condition— $\nu^{f_3} = \dots = \nu^{f_{102}} = (1, 0)$.

Consider the priority profile $f_1 \triangleright_\ell f_2, \dots \triangleright_\ell f_{102}$. In the family-optimal envy-free outcome, only family f_1 is matched to ℓ because ℓ cannot accommodate f_2 and f_2 would envy any family f_j matched to ℓ for $j > 2$. This means that all 99 places for children at ℓ are wasted. In the family-optimal weakly envy-free outcome, in contrast, waste is entirely eliminated as all families except f_2 are matched to ℓ .

Now, consider the priority profile $f_3 \triangleright_\ell \dots \triangleright_\ell f_{102} \triangleright_\ell f_1 \triangleright_\ell f_2$. In the family-optimal envy-free outcome, the place devoted to the rare medical condition is wasted because f_{102} would envy f_1 or f_2 if either were matched to ℓ . However, in the family-optimal weakly envy-free outcome, f_1 can be matched to ℓ , which eliminates waste. \square

Example 2 is merely illustrative; however, our simulations in Section 6 suggest that the efficiency impact of using weak envy-freeness instead of envy-freeness could be significant in practice.

Before looking for a family-optimal weakly envy-free matching in the next section, we clarify the logical relationships among the three priority-respecting solution concepts introduced thus far. Envy-freeness and stability are logically independent: Envy-freeness allows waste which stability precludes; however, stability allows for some waste-eliminating priority violations. Weak envy-freeness is logically independent of stability in the same way as envy-freeness is. In contrast, stability only fails due to the presence of waste or via a priority violation so non-wasteful and envy-free matchings are always stable.²⁸ Perhaps sur-

²⁶Suppose there are two families f_1 and f_2 , one locality ℓ , and one dimension d with $\nu_d^{f_1} = \nu_d^{f_2} = 1$, $\kappa_d^\ell = 2$, $\ell \succ_{f_1} \emptyset$, $\ell \succ_{f_2} \emptyset$, and $f_1 \triangleright_\ell f_2$. The matching $(f_1, \emptyset), (f_2, \ell)$ is weakly envy-free, but not envy-free.

²⁷In reality, there would be at least one more constraint on the total number of refugees but for the sake of illustration, we assume that no constraint other than the two explicitly mentioned are binding.

²⁸Kamada and Kojima (2020) define stability to be the combination of envy-freeness and non-wastefulness, making their definition more restrictive than ours.

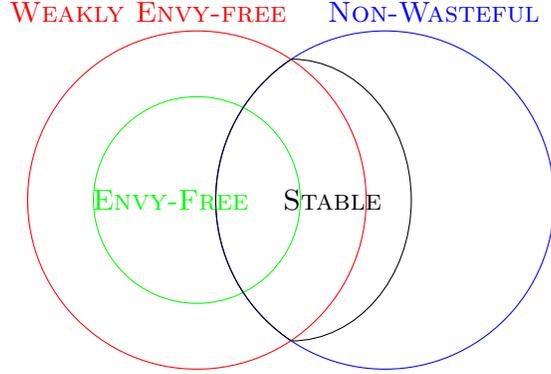


Figure 2: Logical relationships among our solution concepts.

Algorithm 3: KNAPSACK DEFERRED ACCEPTANCE (KDA)

Round $i \geq 1$

Step 1: Every family proposes to its favorite locality that has not permanently rejected it yet.

Step 2: Every locality tentatively accepts a proposing family if the locality can weakly accommodate the family alongside all families with a higher priority that are proposing to or have been permanently rejected by that locality. Otherwise the locality permanently rejects the family.

Step 3: If at least one family has been permanently rejected in Step 2, continue to Round $i + 1$. Otherwise permanently match every family to the locality to which the family is proposing and end.

prisingly, although weak envy-freeness allows for some priority violations, we can establish an analogous relationship among stability, non-wastefulness, and weak envy-freeness.

Proposition 6. *If a matching is non-wasteful and weakly envy-free, then it is stable.*

Proposition 6, combined with the possible nonexistence of a stable matching, formally establishes that a weakly envy-free and non-wasteful matching may not exist. Figure 2 summarizes the relationships among the solution concepts. Note that weak envy-freeness is strictly weaker than envy-freeness even when $|D| = 1$.

5.4 A family-optimal weakly envy-free mechanism

We now prove the existence of the (unique) family-optimal weakly envy-free matching by introducing an algorithm that finds this matching in polynomial time (Algorithm 3). In each

round of our Knapsack Deferred Acceptance (KDA) algorithm (Algorithm 3), all families propose to their favorite localities that have not permanently rejected them yet. A locality ℓ permanently rejects a family f if ℓ cannot weakly accommodate f alongside families with a higher priority from which ℓ has received a proposal (in this round or a previous one). By construction, families that have already proposed to ℓ can only be matched to ℓ or a less-preferred locality, which implies that if ℓ permanently rejects f then f cannot be matched to ℓ in any weakly envy-free matching. If f is not permanently rejected, then f continues to propose to ℓ in the next round. In each round of the KDA algorithm, at least one family is permanently rejected so the algorithm eventually terminates. The KDA algorithm matches each family to its most-preferred locality in any weakly envy-free matching, which yields the following result.

Theorem 5. *The output of the KDA algorithm is the unique family-optimal weakly envy-free matching.*

Since weak envy-freeness allows for some priority violations, the existence of the family-optimal weakly envy-free matching, which is directly implied by Theorem 5, does not follow from the existence of the family-optimal envy-free matching unless $|D| = 1$ (Wu and Roth, 2018; Kamada and Kojima, 2020).

We now illustrate the KDA algorithm and discuss its incentive properties.

Example 3. There are four families, four localities, and one dimension. The preferences, priorities, sizes, and capacities are

$$\begin{array}{cccc} \succ_{f_1}: \ell_2, \ell_1, \dots & \succ_{f_2}: \ell_1, \ell_3, \ell_4, \dots & \succ_{f_3}: \ell_1, \ell_2, \dots & \succ_{f_4}: \ell_1, \ell_3, \dots \\ \triangleright_{\ell_1}: f_1, f_2, f_3, f_4 & \triangleright_{\ell_2}: f_3, f_1, \dots & \triangleright_{\ell_3}: f_4, f_2, \dots & \triangleright_{\ell_4}: \dots \end{array}$$

$$\nu = d_1 \begin{pmatrix} & f_1 & f_2 & f_3 & f_4 \\ 1 & 2 & 1 & 1 & \end{pmatrix} \quad \kappa = d_1 \begin{pmatrix} & \ell_1 & \ell_2 & \ell_3 & \ell_4 \\ 2 & 1 & 2 & 5 & \end{pmatrix}.$$

It can be verified that this market does not have a stable matching. Its family-optimal weakly envy-free matching is

$$\begin{pmatrix} f_1 & f_2 & f_3 & f_4 \\ \ell_1 & \ell_4 & \ell_2 & \ell_3 \end{pmatrix}.$$

The family-optimal weakly envy-free matching can be calculated using the KDA algorithm, which is displayed in Table 1.²⁹

In Round 1, all families propose to their favorite localities. Locality ℓ_1 has two units of capacity and receives proposals from families f_2 , f_3 , and f_4 . As family f_2 has the highest

²⁹As there is only one dimension, this is also the family-optimal envy-free matching (see Corollary 3).

Round 1	Round 2	Round 3	Round 4	Round 5
$f_1 \rightarrow \ell_2 \checkmark$	$f_1 \rightarrow \ell_2 \times$	$f_1 \rightarrow \ell_1 \checkmark$	$f_1 \rightarrow \ell_1 \checkmark$	$f_1 \rightarrow \ell_1 \checkmark$
$f_2 \rightarrow \ell_1 \checkmark$	$f_2 \rightarrow \ell_1 \checkmark$	$f_2 \rightarrow \ell_1 \times$	$f_2 \rightarrow \ell_3 \times$	$f_2 \rightarrow \ell_4 \checkmark$
$f_3 \rightarrow \ell_1 \times$	$f_3 \rightarrow \ell_2 \checkmark$			
$f_4 \rightarrow \ell_1 \times$	$f_4 \rightarrow \ell_3 \checkmark$			

Table 1: KDA algorithm applied to Example 3.

priority and a size of two, f_2 is tentatively accepted while the other two families are permanently rejected. In Round 2, families f_3 and f_4 propose to their second-choice localities ℓ_2 and ℓ_3 respectively. Families f_1 and f_3 compete for the single unit of capacity in locality ℓ_2 . Family f_3 is tentatively accepted as f_3 has a higher priority than f_1 ; therefore f_1 is permanently rejected. In Round 3, family f_1 proposes to its second-choice locality ℓ_1 and competes with f_2 . As locality ℓ_2 has a capacity of two but families f_1 and f_2 joint size is three, only the family with the highest priority (i.e., f_1) is tentatively accepted. Hence, locality ℓ_1 permanently rejects family f_2 . In Round 4, family f_2 proposes to locality ℓ_3 and is also permanently rejected since f_4 has a higher priority than f_2 . In Round 5, f_2 finally proposes to ℓ_4 . There are no permanent rejections; therefore, all families are now permanently matched and the algorithm ends. \square

Example 3 sheds light on two important aspects of the KDA algorithm. First, Example 3 shows how some units of capacity can remain unused. In Round 1, families f_3 and f_4 are permanently rejected because f_2 is taking both units of locality ℓ_1 's capacity. In Round 3, however, family f_2 is permanently rejected by locality ℓ_1 because ℓ_1 receives a proposal from f_1 . As family f_1 's size is one, locality ℓ_1 's second unit of capacity remains unused although this unit could be used by families f_3 or f_4 . In the case when there is only one dimension, the maximum number of units that remain unused in each locality is equal to the largest size of any family minus one. In our example, the largest size is two (family f_2 's size) so at most one unit of capacity may end up unused at each locality. If there are multiple dimensions, this bound remains valid for at least one dimension in each locality. The number of unused units for other dimensions may be larger.

Second, Example 3 also reveals that the KDA algorithm is not strategy-proof. To see this, suppose that family f_2 reports locality ℓ_3 to be its first choice. Then, all families are tentatively accepted in the first round since locality ℓ_1 has two units of capacity and families f_3 and f_4 each require one unit; f_1 and f_2 are the only families that propose to ℓ_2 and to ℓ_3 , respectively. The KDA outcome is

$$\begin{pmatrix} f_1 & f_2 & f_3 & f_4 \\ \ell_2 & \ell_3 & \ell_1 & \ell_1 \end{pmatrix}.$$

Algorithm 4: THRESHOLD KNAPSACK DEFERRED ACCEPTANCE (TKDA)

Round $i \geq 1$

Step 1: Every family f proposes to its favorite locality ℓ that has not permanently rejected f yet.

Step 2: Every locality ℓ permanently rejects any proposing family f if f 's priority rank among all families that are proposing to ℓ is strictly greater than f 's threshold at ℓ (calculated by Algorithm 5).

Step 3: If at least one family has been permanently rejected in Step 2, continue to Round $i + 1$. Otherwise permanently match every family to the locality to which the family is proposing and end.

The following result is an immediate consequence of this observation.

Proposition 7. *There is no strategy-proof and family-optimal weakly envy-free mechanism.*

The driving force behind Proposition 7 is that the acceptance rule (choice function) of localities induced by the KDA algorithm does not satisfy the cardinal monotonicity condition (Alkan, 2002; Alkan and Gale, 2003; Fleiner, 2003; Hatfield and Milgrom, 2005).³⁰ Cardinal monotonicity requires that the number of tentatively accepted families grows monotonically with the number of proposing families. Cardinal monotonicity is crucial for designing strategy-proof matching mechanisms (Hatfield and Milgrom, 2005). Let us see how cardinal monotonicity is violated in Example 3. If families f_3 and f_4 were to propose to locality ℓ_1 , they would be tentatively accepted. However, if families f_3 , f_4 , and f_2 were to propose to locality ℓ_1 together, then only family f_2 would be tentatively accepted. Subsequent to family f_3 's permanent rejection, there is a rejection chain that leads to family f_2 's being permanently rejected by locality ℓ_1 . At the same time, having been permanently rejected by locality ℓ_1 , family f_4 ends up competing with f_2 at ℓ_3 . Therefore, f_2 is not only permanently rejected by its first-choice locality, but also faces more competition in its second-choice locality. As a result, f_2 has the incentive to misreport its preferences and propose directly to ℓ_3 in order avoid being adversely affected by the two rejection chains. If cardinal monotonicity is satisfied, then f_2 's proposal can create at most one rejection chain.

Algorithm 5: THRESHOLD CALCULATOR

For every locality $\ell \in L$, let Π_ℓ be the set of families that are currently proposing to ℓ .

Step 1: For every family $f \in F$ and every locality $\ell \in L$, calculate the *temporary threshold* of f at ℓ , denoted by $\tilde{\theta}_\ell^f$, as follows:

- If ℓ can weakly accommodate f alongside \widehat{F}_ℓ^f , let $\tilde{\theta}_\ell^f = \infty$.
- If ℓ cannot weakly accommodate f alongside $\Pi_\ell \cap \widehat{F}_\ell^f$, let $\tilde{\theta}_\ell^f = 0$.
- Otherwise, find the unique $n \in \mathbb{Z}_{>0}$, such that
 - (i) ℓ can weakly accommodate f alongside all sets of families $G \subseteq \widehat{F}_\ell^f$ such that $|G| = n - 1$ and $(\Pi_\ell \cap \widehat{F}_\ell^f) \subseteq G$, and
 - (ii) ℓ cannot weakly accommodate f alongside a set of families $G' \subseteq \widehat{F}_\ell^f$ such that $|G'| = n$ and $(\Pi_\ell \cap \widehat{F}_\ell^f) \subseteq G'$;
 and let $\tilde{\theta}_\ell^f = n$.

Step 2: For every family $f \in F$ and every locality $\ell \in L$, calculate the *threshold* of f at ℓ , θ_ℓ^f , as follows:

- If $\tilde{\theta}_\ell^f = \infty$, let $\theta_\ell^f = \infty$.
- Otherwise, let $\theta_\ell^f = \min_{g \in \widehat{F}_\ell^f \cup \{f\}} \tilde{\theta}_\ell^g$.

5.5 A weakly envy-free and strategy-proof mechanism

In this section, we introduce the Threshold Knapsack Deferred Acceptance (TKDA) algorithm (Algorithm 4). We show that the TKDA mechanism is weakly envy-free and strategy-proof. A direct consequence of Proposition 7 is that the TKDA algorithm will not always produce the family-optimal weakly envy-free matching, but, in Online Appendix A.1 we show that both the KDA and TKDA algorithms are guaranteed to match a minimum number of families (Proposition 9).

The TKDA algorithm follows the structure of the KDA algorithm. In each round, every family proposes to its most preferred locality from which the family has not yet been permanently rejected. Localities tentatively accept some proposals and permanently reject others. This process continues until all proposals are tentatively accepted. In that round, all families are permanently matched to the last locality to which they proposed and the algorithm ends.

The key part of the TKDA algorithm is the choice rule that decides whether a proposal is tentatively accepted or permanently rejected (Algorithm 5). In order to ensure that the TKDA algorithm is strategy-proof, the choice rule of localities needs to satisfy cardinal monotonicity. That is, for each locality, the choice rule ensures that as the set of proposing families expands, the number of families that are tentatively accepted by the locality weakly increases. The choice rule could therefore result in fewer families being tentatively accepted than the locality can weakly accommodate. Consequently, the matching produced by the TKDA algorithm may not be family-optimal in general and additional capacity may remain unused as a result.

At a high level, the choice rule works as follows (see Step 2 of Algorithm 4). In every round, every locality assigns every family (not just those that are proposing) a *threshold*. For every proposing family, the locality then compares the family’s threshold with its priority rank among proposing families. The family is tentatively accepted so long as its priority rank does not exceed its threshold, otherwise the family is permanently rejected.

The details of the threshold calculation are more involved (Algorithm 5). First, every locality ℓ assigns every family f a *temporary* threshold (Step 1 of Algorithm 5). If locality ℓ can weakly accommodate f alongside all higher-priority families \widehat{F}_ℓ^f (whether or not they are proposing), f ’s temporary threshold at ℓ is ∞ (ensuring that ℓ will tentatively accept a proposal from f no matter which other families are proposing). If ℓ cannot weakly accommodate f alongside all higher-priority families \widehat{F}_ℓ^f that are currently proposing to ℓ , f ’s temporary threshold at ℓ is 0 (ensuring that ℓ will permanently reject any future proposal

³⁰In other contexts, cardinal monotonicity has been referred to as “size monotonicity” (Alkan and Gale, 2003) and the “Law of Aggregate Demand” (Hatfield and Milgrom, 2005).

	Round 1	Round 2	Round 3	Round 4	Round 5
$f_1 \rightarrow$	ℓ_2 [1] ✓	ℓ_2 [0] ✗	ℓ_1 [∞] ✓	ℓ_1 [∞] ✓	ℓ_1 [∞] ✓
$f_2 \rightarrow$	ℓ_1 [1] ✓	ℓ_1 [1] ✓	ℓ_1 [0] ✗	ℓ_3 [0] ✗	ℓ_4 [∞] ✓
$f_3 \rightarrow$	ℓ_1 [0] ✗	ℓ_2 [∞] ✓			
$f_4 \rightarrow$	ℓ_1 [0] ✗	ℓ_3 [∞] ✓			

Table 2: TKDA algorithm applied to the market from Example 3 with true reporting. Thresholds are in square brackets.

from f since matching f to ℓ would create strong envy from at least one of the families that are currently proposing). Otherwise, f 's threshold is a positive integer, calculated as follows. We consider all sets of higher-priority families that contain those higher-priority families that are currently proposing to ℓ . Then, we assign f a temporary threshold n for which (i) ℓ can weakly accommodate f alongside all such sets containing $n - 1$ families but (ii) ℓ cannot weakly accommodate f alongside at least one such set of n families. The threshold of family f at ℓ is the smallest temporary threshold among its own and the temporary thresholds of all families with a higher priority than f at ℓ .³¹ The calculation of thresholds can be done in polynomial time, which ensures that the TKDA algorithm is practical even for large markets.³²

Theorem 6. *The TKDA mechanism is strategy-proof and weakly envy-free.*

In Example 3, the TKDA algorithm follows the same deferred acceptance procedure and produces the same matching as the KDA algorithm. The TKDA algorithm for Example 3 is displayed in Table 2. In Round 1, families f_2 , f_3 , and f_4 propose to locality ℓ_1 . The thresholds of families f_3 and f_4 are 0 since locality ℓ_1 cannot weakly accommodate either one of those families alongside f_2 . Hence, families f_3 and f_4 are permanently rejected by locality ℓ_1 . In Rounds 2-4, families f_1 and f_2 are assigned a threshold of 0 (at locality ℓ_2 and at localities ℓ_1 and ℓ_3 respectively) and are permanently rejected by these localities. The algorithm ends in Round 5, in which no permanent rejections occur because every family's threshold is infinity.³³

³¹This step precludes any priority violation among families with a finite threshold, which is essential to preserve cardinal monotonicity; however, the matching produced by the TKDA algorithm may still not be envy-free as families with an infinite threshold can still be envied by higher-priority families with finite thresholds.

³²Take each dimension for which f 's size is at least one and order the families with a higher priority from largest to smallest size for that dimension. Starting from the largest family, add one family at a time until the families' total size for that dimension (including f 's) exceeds the capacity. A number of families is obtained in this way for each dimension, n is the minimum among these numbers.

³³Example 3 is simple and aimed at illustrating the main difference between KDA and TKDA. For a full worked out example of the TKDA algorithm, see Online Appendix B.3.

Round 1				Round 2				Round 3			
f_1	\rightarrow	ℓ_2	[1] ✓	f_1	\rightarrow	ℓ_2	[1] ✓	f_1	\rightarrow	ℓ_2	[1] ✓
f_2	\rightarrow	ℓ_3	[1] ✓	f_2	\rightarrow	ℓ_3	[0] ✗	f_2	\rightarrow	ℓ_4	[∞] ✓
f_3	\rightarrow	ℓ_1	[1] ✓	f_3	\rightarrow	ℓ_1	[1] ✓	f_3	\rightarrow	ℓ_1	[1] ✓
f_4	\rightarrow	ℓ_1	[1] ✗	f_4	\rightarrow	ℓ_3	[∞] ✓	f_4	\rightarrow	ℓ_3	[∞] ✓

Table 3: TKDA algorithm applied to the market from Example 3 with a misreport. Thresholds are in square brackets.

Even though both algorithms produce the same matching, the TKDA algorithm removes family f_2 's incentive to misreport its preferences that f_2 has in the KDA algorithm. Suppose that family f_2 manipulates its preferences report to $\succ_{f_2}: \ell_3, \ell_4, \dots$. Under the KDA algorithm, such a manipulation allowed family f_2 to be matched to locality ℓ_3 instead of ℓ_4 . This is no longer possible in the TKDA algorithm as we illustrate in Table 3.

In Round 1, following family f_2 's manipulation, families f_3 and f_4 propose to ℓ_1 . Locality ℓ_1 can weakly accommodate both families but cannot weakly accommodate either one alongside f_2 ; therefore the threshold of both families is 1. Locality ℓ_1 therefore tentatively accepts the proposing family with the highest priority (f_3) but permanently rejects the one with the second-highest priority (f_4). This is the key step in which TKDA differs from KDA: locality ℓ_1 permanently rejects family f_4 in Round 1 whether or not it receives a proposal from f_2 . In Round 2, families f_2 and f_4 propose to locality ℓ_3 . As family f_4 has the highest priority for ℓ_3 , its threshold is ∞ and f_4 is tentatively accepted by ℓ_3 . In contrast, locality ℓ_3 cannot weakly accommodate family f_2 alongside f_4 ; therefore family f_2 's threshold is 0 and f_2 is permanently rejected by ℓ_3 . In Round 3, family f_2 proposes to locality ℓ_4 and the algorithm ends as no family is permanently rejected. Note that family f_2 's manipulation has benefited f_1 and f_3 , which are matched to their first- rather than second-choice localities. However, family f_2 has not managed to benefit from its own manipulation and has remained matched to locality ℓ_4 .

One might be concerned that Theorem 6 does not provide a lower bound on the efficiency of the TKDA mechanism. In Online Appendix A.1, we derive a lower bound for the efficiency of the TKDA algorithm (Proposition 9). Finally, we provide a way to improve the efficiency of the TKDA mechanism without affecting its properties (Algorithm 7). Before the TKDA algorithm is run, we identify family-locality pairs that will necessarily be matched together. For every such pair, the family *clinches* the locality and is no longer considered by less-preferred localities. In the TKDA algorithm, clinching may allow us to raise the thresholds of the remaining families. As a result, the matching found by the TKDA with Clinching (TKDAC) mechanism weakly Pareto dominates the matching found by the TKDA mechanism (see Online Appendix A.2 for details).

5.6 Properties of mechanisms in a one-dimensional setting

As we have showed in Section 5.3, even when $|D| = 1$, there exist matchings that are weakly envy-free but not envy-free (footnote 26). However, it turns out that neither the KDA nor the TKDA algorithm will ever produce such a matching.

Proposition 8. *If $|D| = 1$, then the KDA and TKDA mechanisms are envy-free.*

Since the KDA algorithm produces the family-optimal weakly envy-free matching (Theorem 5), we know from Proposition 8 that this matching must also be envy-free when $|D| = 1$. As all envy-free matchings are weakly envy-free and there exists a unique family-optimal envy-free matching (Proposition 5), the KDA algorithm produces the family-optimal envy-free matching.

Corollary 3. *If $|D| = 1$, then the family-optimal envy-free matching and the family-optimal weakly envy-free matching coincide.*

6 Data and Simulations

Our theoretical results leave important empirical questions open about how our mechanisms perform in practice. To get those questions, here we simulate our mechanisms using data from HIAS. We begin with the analysis of one-dimensional environment in which each locality only has a constraint on the number of refugees—this is closest to HIAS’s current practice. We then consider how our matching mechanisms perform in a three-dimensional setting.

Data. We use anonymized data on refugees resettled by HIAS from October 2016 to October 2017. There are 839 refugees partitioned into 329 families that were matched across 20 localities. To preserve anonymity, our data contain only the number of senior, (working-age) adult, and child family members. Following Nguyen et al. (2019) and Trapp et al. (2020), we conservatively treat each locality’s quota as the total number of refugees that HIAS actually resettled in that locality in 2017. The total number of refugees resettled in a locality varied between 5 and 99.³⁴ There were 332 children, 498 adults, and 9 seniors in our data. The largest family had 8 members and there were 157 families with only 1 member. We use expected employment estimates from Trapp et al. (2020) to determine priorities: for each family we add the employment likelihoods of all working-age adults to determine the expected employment *weight* of each family and localities rank families from highest to lowest

³⁴The annual quotas were abruptly adjusted following a change in the Presidential administration in January 2017.

		Preferences vs. Priorities	
		Correlated	Uncorrelated
Preferences	Correlated	Type 4	Type 1
	Uncorrelated	Type 3	Type 2

Table 4: Preference types for simulations.

expected employment weight. Since families that have high employment weights in some localities are likely to have high employment weights in other localities, induced priorities are naturally correlated. We make an exception for the KTTCE mechanism, where localities give a higher priority to families in their endowment (but otherwise rank families in decreasing likelihood of employment).

Simulated preferences. As HIAS does not currently collect preferences from refugees, we have to simulate preferences for our analysis. We assume that the utility of family f associated with being assigned to locality ℓ takes the form

$$U_{f\ell} = \delta V_{f\ell} + \beta Y_{\ell} + \gamma E_{f\ell},$$

where $V_{f\ell}$ is the (normalized to $[0,1]$) employment weight of family f in locality ℓ , Y_{ℓ} is a locality-specific term, and the pure noise term $E_{f\ell}$. Terms Y_{ℓ} (for each locality) and $E_{f\ell}$ (for each family-locality pair) are drawn independently from $U(0,1)$ in each simulation round. We consider four types of preferences, which are summarized in Table 4.

- Type 1: $\beta = 1, \delta = \gamma = 0$. In each simulation round, we draw a value for each locality that determines all of the family preferences. Therefore, preferences are perfectly correlated among families but uncorrelated with the priorities of localities.
- Type 2: $\gamma = 1, \beta = \delta = 0$. In each simulation round, we draw a value for each family-locality pair. Therefore, preferences are uncorrelated among families and uncorrelated with priorities.
- Type 3: $\delta = \gamma = 1, \beta = 0$. In each simulation round, family preferences are a sum of their employment weight and a noise term. Therefore, preferences are uncorrelated among families but correlated with priorities.
- Type 4: $\delta = \beta = 1, \gamma = 0$. In each simulation round, family preferences are a sum of their employment weight and a locality-specific term. Therefore, preferences are correlated among families and correlated with priorities.

Preference type	Type 1	Type 2	Type 3	Type 4
Families made better off	6.57	39.45	39.05	21.54

Table 5: Number of families made better off by the KTTCE mechanism (one dimension). Averages over 100 simulation rounds.

For each preference type, in each of the 100 simulation rounds, we computed the outcome of KDA, TKDA, KTTC, and KTTCE (where the endowment is the employment-maximizing allocation) algorithms.³⁵ In this section, we report means for various statistics taken over all the simulation rounds.

6.1 Performance of the KTTCE mechanism.

We begin by analyzing the performance of the KTTCE mechanism, as we see that mechanism as representing a first step towards incorporating preferences in refugee resettlement processes. We take the employment-maximizing matching as the endowment; that is, every family is endowed with the locality to which it is matched under the matching that maximizes total refugee employment. Each locality prioritizes the families that are in its endowment over those that are not and, within each tier, ranks family from highest to lowest expected employment. Note that in this setting, Theorem 4 does not apply because priorities are not lexicographic. Therefore, our theoretical results do not guarantee KTTCE that would find any Pareto improvements from the endowment.

The results are presented in Table 5. The number of families that the KTTCE mechanism makes better off is small for Type 1 preferences, with just under 7 families helped on average. The reason is that Type 1 families cannot trade with each other as they have identical preferences; therefore the only improvements that the KTTCE mechanism finds are closed chains, in which a family takes advantage of unused capacity in a more preferred locality. The KTTCE mechanism performs best with Type 2 and Type 3 preferences because family preferences are uncorrelated, which creates the possibility of many mutually beneficial trades. Overall, the simulation results in Table 5 are encouraging—they suggest that the KTTCE mechanism would find at least some of the Pareto improvements in practice even when it is not guaranteed to.

A part of the design of the KTTCE mechanism that we leave open is the order in which families are picked in the Rejection Stage. In our main simulations, that order is determined randomly (with an equal probability for each possible order). We also test two alternative rules: ordering families from largest to smallest and from smallest to largest

³⁵Our code is available at <https://github.com/nhhai196/Refugee-Resettlement>.

Preference type	Type 1	Type 2	Type 3	Type 4
Largest-to-smallest family	6.57	39.67	40.56	21.87
Smallest-to-largest family	4.89	10.61	11.85	8.11

Table 6: Number of families made better off for different orderings in the Rejection Stage of the KTTCE (one dimension). Averages over 100 simulation rounds.

(same-size families continue to be ordered randomly). The results of the two alternative family orderings in the Rejection Stage are displayed in Table 6. We see that the largest-to-smallest ordering is more efficient than the smallest-to-largest ordering. Intuitively, larger families are more likely to cause a cycle not to be feasible (see Algorithm 2) as the locality at which they are pointing may not be able to accommodate them. Therefore, by rejecting larger families first, the algorithm needs to reject fewer families before finding a feasible cycle. Compared to a random order, one would then naturally expect rejecting larger families first to have a positive impact on efficiency and rejecting smaller families first to have a negative one. Why is the efficiency gain of the largest-to-smallest ordering small relative to a random ordering? The reason is that families that can be rejected in the Rejection Stage tend to be large, meaning that, when families are ordered randomly, it is likely that larger families will be picked first and the algorithm will find a feasible cycle before picking one of the few small families. Consequently, a random ordering does not result in large efficiency losses compared to the largest-to-smallest ordering.

6.2 Comparison of the KTTC, KDA, and TKDA mechanisms

We next turn to our three remaining mechanisms: KTTC, KDA, and TKDA. We compare those mechanisms in terms of strong envy and efficiency. The results are presented in Table 7 and Figure 3.

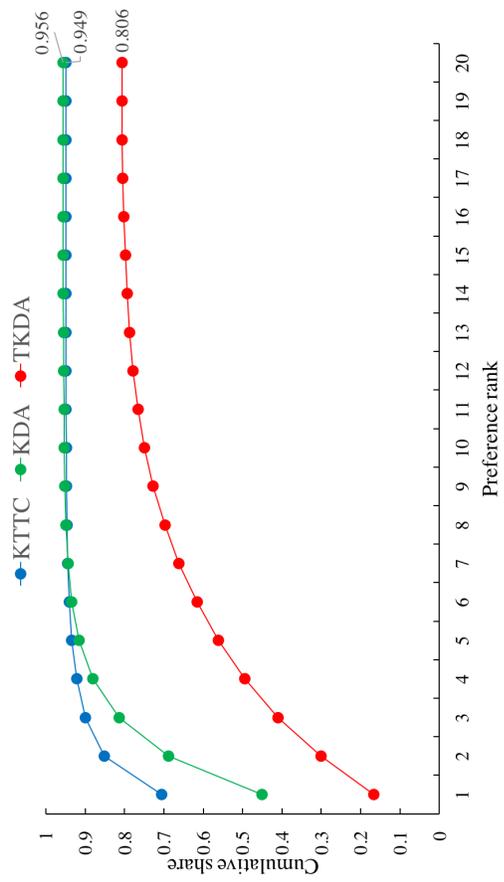
To count the number of strong envy violations in each simulation round, we count the number of family pairs (f_1, f_2) such that f_1 strongly envies f_2 .³⁶ In line with Theorems 5 and 6, the number of strong envy violations for the KDA and TKDA mechanisms is 0 because they are weakly envy-free. For the KTTC mechanism, the number of strong envy violations is very small for Type 1 preferences (so KTTC mechanism coincides with a serial dictatorship). The reason is that the preferences of all families are identical so strong envy violations only occur when a larger family cannot be accommodated in a locality but a smaller family with a lower priority can.³⁷ The number of strong envy violations increases for Type 2 and Type 3

³⁶As the ordering of the two families matter, there are $329^2 = 108241$ such pairs.

³⁷For example, suppose that locality ℓ has a capacity 2 with $|D| = 1$. There are two families f_1 and f_2 who want to be matched to ℓ , such that f_1 has a higher priority but is of size 3 while f_2 is of size 2. The

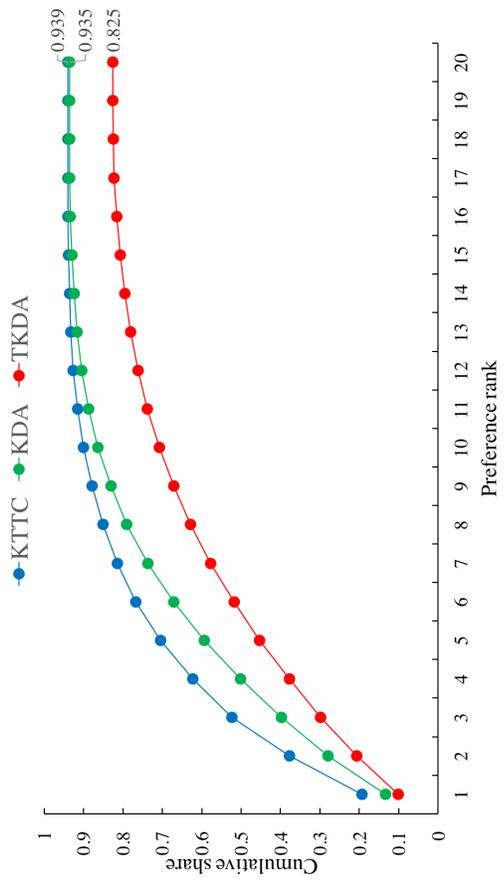
Preference type	Type 1	Type 2	Type 3	Type 4
Strong envy violations				
KTTC	39	762	730	3639
KDA	0	0	0	0
TKDA	0	0	0	0
Number of matched families				
KTTC	308	312	312	309
KDA	305	315	315	308
TKDA	276	265	265	271
Fraction of unfilled capacity				
KTTC	7.0%	4.8%	5.6%	7.2%
KDA	10.5%	7.7%	7.7%	10.0%
TKDA	18.6%	20.9%	20.7%	19.6%

Table 7: Outcomes of KTTC, KDA, and TKDA algorithms (one dimension). Averages over 100 simulation rounds. Numbers in the top and middle panels are rounded to nearest integer; numbers in the bottom panel are rounded to 1 d.p.



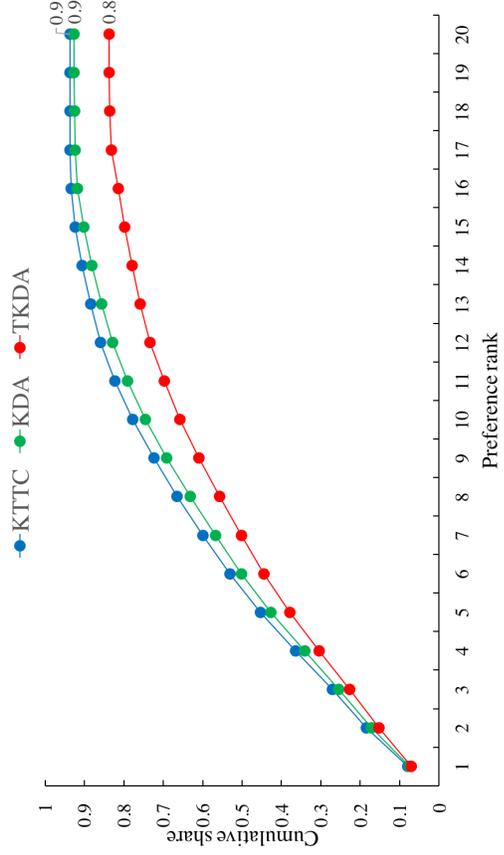
(a) Type 1

(b) Type 2



(c) Type 3

Figure 3: Preference distributions for outcomes under different matching mechanisms (one dimension). Labelled numbers: fractions of matched families). Averages over 100 simulation rounds.



(d) Type 4

preferences because heterogeneous preferences means that more priority trades are possible, which may create strong envy violations. Type 4 preferences capture an intermediate “worst-case” scenario for strong envy: as preferences are correlated there are not many opportunities for trade, however, since priorities and preferences are correlated many families that do not get their top-choice localities are likely to envy lower-priority families that ended up in these localities.

Figure 3 illustrates the efficiency of the three mechanisms. Recall that Theorems 5 and 6 imply that every family is weakly better off under the KDA mechanism than under the TKDA mechanism. Figure 3 shows that the KDA mechanism is substantially more efficient than the TKDA mechanism, especially when preferences are uncorrelated (Type 2 and Type 3). In addition, the two bottom panels in Table 7 show that the KDA mechanism more families and leaves less unfilled quota than the TKDA mechanism. Overall, our simulations suggest that the strategy-proofness of TKDA could come a significant efficiency cost.

Our theoretical results are silent about whether the efficiency ranking between KTTC and KDA mechanisms and indeed our simulations suggest that, depending on the preference type, some families can be better off under either mechanism. However, as the KTTC mechanism is Pareto-efficient, we should expect it to be more efficient “on average” than the Pareto-inefficient KDA mechanism. Figure 3 confirms this intuition for all preference types. However, Figure 3 also suggests that the difference between the KTTC and KDA mechanisms is small compared to the difference between the KDA and TKDA mechanisms. Consequently, our simulations indicate that in practice the efficiency cost of imposing weak envy-freeness may not be particularly large.

6.3 Comparison of mechanisms under three-dimensional constraints

The previous section considered a one-dimensional setting in which the only constraint is on the number of refugees that each locality can resettle. The one-dimensional setting is in line with current practice at HIAS and other US resettlement agencies. However, as we mention in Section 3.1, the possibility of accounting for multidimensional constraints can offer resettlement agencies greater flexibility in resettlement matching. In this section, we thus test our mechanisms in an environment with multidimensional constraints. Our data contains information on the number of children, adults, and seniors in each family, which we use to construct our three constraint dimensions.³⁸ That is, each family’s size is a three-

KTTC mechanism will match f_2 to ℓ so f_3 will strongly envy f_2 while the KDA and TKDA mechanisms will not match either family to ℓ .

³⁸We imagine here that localities could limit the total number of refugees of different ages because of age-relevant local constraints such as kindergarten places, employment training, or elderly care.

Preference type	Type 1	Type 2	Type 3	Type 4
Families made better off	22.39	81.62	83.55	46.73

Table 8: Number of families made better off by the KTTCE mechanism (three dimensions). Averages over 100 simulation rounds.

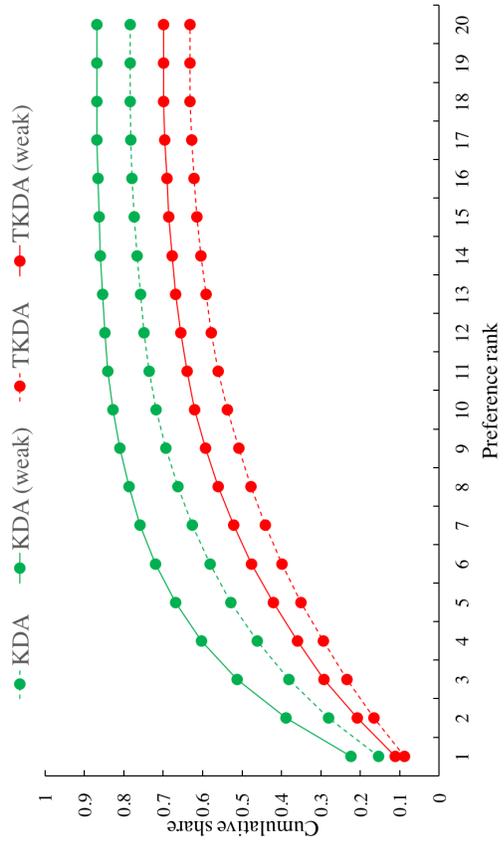
dimensional vector consisting of the number of children, adults, and senior family members and the capacity of each locality is a three-dimensional vector consisting of the number of children, adults, and seniors that the locality resettled in 2017. The three-dimensional setting is strictly more constrained than the one-dimensional setting.³⁹ However, as Table 13 and Figure 10 in Online Appendix C show, the results for the KTTC, KDA, and TKDA algorithms for the three-dimensional setting are qualitatively similar to the results in the one-dimensional setting.

Table 8 summarizes the performance of the KTTCE mechanism in the three-dimensional environment.⁴⁰ The number of families made better off on average is more than double than in the one-dimensional environment across all preference types. This result may look counterintuitive at first sight because one might assume that tougher constraints make it harder for the KTTCE mechanism to find improvements. However, the integer optimization program that determines the employment-maximizing endowment matching also faces more constraints and therefore leaves a large number of families unmatched. As a result, the KTTCE mechanism is able to take advantage of the unfilled capacity to make more families better off.

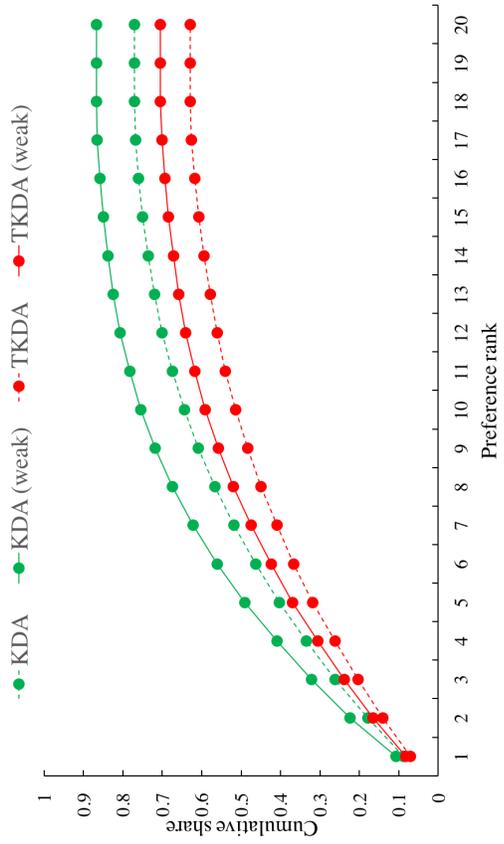
Finally, we assess the efficiency gain associated with using our solution concept of weak envy-freeness as opposed to envy-freeness. We modify the KDA algorithm so that it produces the family-optimal *envy-free matching* and the TKDA algorithm so that it produces an envy-free matching (that is, by only replacing “weak accommodation” with “accommodation” in Algorithms 3 and 5). With one dimension, the KDA and TKDA mechanisms are envy-free (Proposition 8) but in the presence of multidimensional constraints, weak envy-freeness is strictly less demanding. Figure 4 shows that the impact of using weak envy-freeness on efficiency is substantial both for the KDA and the TKDA algorithms.

³⁹By construction, a matching that satisfies the three constraints also satisfies the original constraint on the number of refugees. Therefore, the three-dimensional setting is equivalent to a four-dimensional setting in which the dimensions are: number of refugees, number of children, number of adults, and number of seniors.

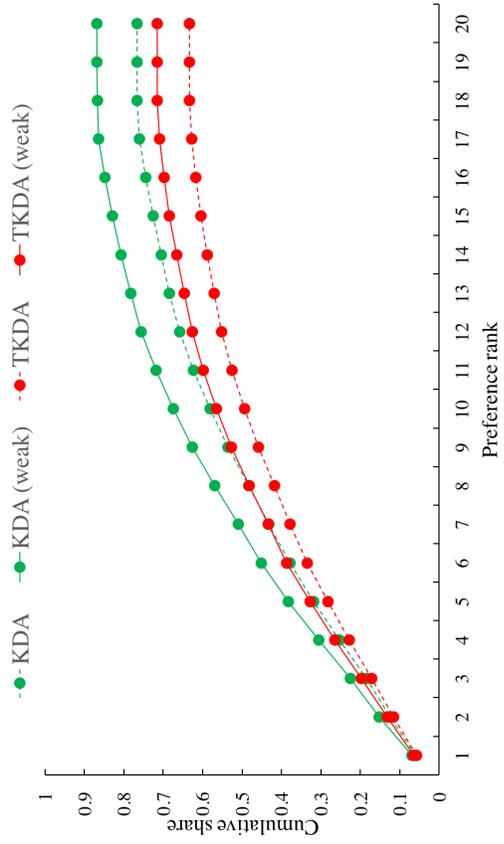
⁴⁰Sizes are not monotonic since a family can have more adults but fewer children than another family.



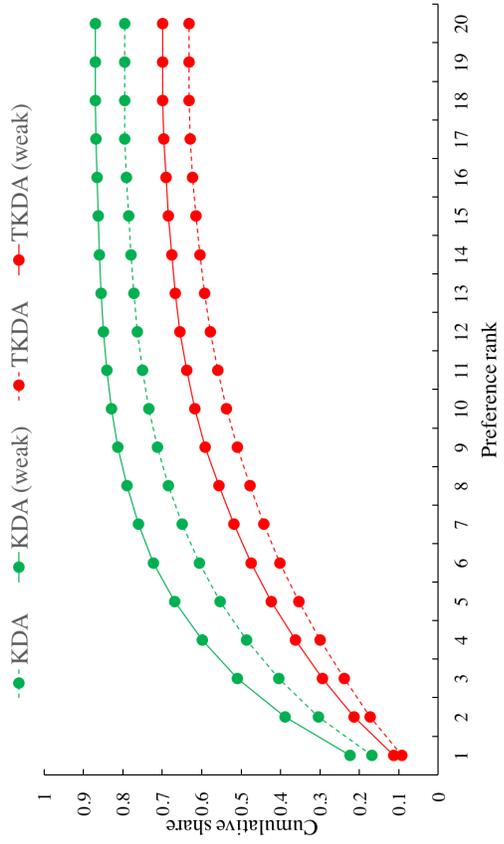
(a) Type 1



(b) Type 2



(c) Type 3



(d) Type 4

Figure 4: Preference distributions for KDA and TKDA algorithms when using accommodation vs. weak accommodation (three dimensions).

7 Conclusion

Refugee resettlement presents a real opportunity for marketplace design: policymakers and resettlement agencies are already working with market design experts to improve matching outcomes in ways that—if we do this work well—stand to improve the lives of millions of disenfranchised people worldwide (Andersson, 2017; Jones and Teytelboym, 2017b; Kominers et al., 2017; Roth, 2018). Recent efforts have focused on maximizing short-run employment outcomes; we show how to take this work further by integrating refugees’ preferences and localities’ priorities into the assignment process. As we have highlighted, the trade-off among respecting priorities of localities, maximizing refugee welfare, and strategy-proofness is exacerbated by the presence of (possibly multidimensional) knapsack constraints in refugee resettlement matching.

We proposed four matching mechanisms for refugee resettlement: KTTC, which modifies the classical Top Trading Cycles (TTC) algorithm to account for knapsack constraints—and, like TTC, is strategy-proof and finds a Pareto efficient outcome; KTTCE, which further extends KTTC to incorporate endowments; KDA, a variant of deferred acceptance that is weakly envy-free and family-optimal in our context; and TKDA, which trades off some efficiency relative to KDA in exchange for strategy-proofness. We have also simulated our mechanisms with data from a resettlement agency to explore how our mechanisms might perform in practice.

Incorporating refugees’ preferences into refugee matching will serve not just to improve the quality of assignment outcomes, but also to give refugees a larger role in the resettlement process. Moreover, collecting information about refugees’ preferences will enable us to better understand what constitutes a high-quality refugee–locality match. Locality priorities are similarly important: if we want localities to be willing to host a large number of refugees, then we must do what we can to respond to their desires and constraints. The hope is that a well-designed resettlement matching system will increase localities’ overall willingness to host refugees, boosting resettlement overall.

Although our simulations already illuminate the key trade-offs between our mechanisms, it will be essential to test the mechanisms using actual preference data. Additionally, we will need to think carefully about how static matching frameworks should account for the dynamic nature of refugee arrival and resettlement (Andersson et al., 2018; Caspari, 2019).

Refugee resettlement matching provides an opportunity for market design to make a real difference into people’s lives. While the long-term success of resettlement programs depends on many factors, improving match quality will make sure that as many refugees as possible get off to a good start.

References

- Abdulkadirođlu, A. (2005). College admissions with affirmative action. *International Journal of Game Theory* 33(4), 535–549.
- Abdulkadirođlu, A., P. A. Pathak, and A. E. Roth (2009). Strategyproofness versus efficiency in matching with indifferences: Redesigning the NYC high school match. *American Economic Review* 99(5), 1954–1978.
- Abdulkadirođlu, A. and T. Sönmez (2003). School choice: A mechanism design approach. *American Economic Review* 93(3), 729–747.
- Alkan, A. (2002). A class of multipartner matching markets with a strong lattice structure. *Economic Theory* 19, 737–746.
- Alkan, A. and D. Gale (2003). Stable schedule matching under revealed preference. *Journal of Economic Theory* 112(2), 289–306.
- Andersson, T. (2017). Refugee matching as a market design application. Working Paper 2017:16, Lund University.
- Andersson, T. and L. Ehlers (2020). Assigning refugees to landlords in Sweden: Efficient stable maximum matchings. *Scandinavian Journal of Economics* 122(3), 937–965.
- Andersson, T., L. Ehlers, and A. Martinello (2018, October). Dynamic refugee matching. Working Paper 2018:7, Lund University.
- Åslund, O., P.-A. Edin, P. Fredriksson, and H. Grönqvist (2011). Peers, neighborhoods, and immigrant student achievement: Evidence from a placement policy. *American Economic Journal: Applied Economics* 3(2), 67–95.
- Åslund, O. and P. Fredriksson (2009). Peer effects in welfare dependence quasi-experimental evidence. *Journal of Human Resources* 44(3), 798–825.
- Åslund, O., J. Östh, and Y. Zenou (2010). How important is access to jobs? Old question, improved answer. *Journal of Economic Geography* 10(3), 389–422.
- Åslund, O. and D.-O. Rooth (2007). Do when and where matter? Initial labour market conditions and immigrant earnings. *The Economic Journal* 117(518), 422–448.
- Aygün, O. and I. Bó (2020, forthcoming). College admission with multidimensional privileges: The Brazilian affirmative action case. *American Economic Journal: Microeconomics*.

- Aziz, H., J. Chen, S. Gaspers, and Z. Sun (2018). Stability and pareto optimality in refugee allocation matchings. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, pp. 964–972. International Foundation for Autonomous Agents and Multiagent Systems.
- Balinski, M. and T. Sönmez (1999). A tale of two mechanisms: Student placement. *Journal of Economic Theory* 84(1), 73–94.
- 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., 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 E. McDermid (2014). Matching with sizes (or scheduling with processing set restrictions). *Discrete Applied Mathematics* 164, 61–67.
- Budish, E. (2011). The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy* 119(6), 1061–1103.
- Caspari, G. (2019, November). An alternative approach to asylum assignment. Mimeo.
- Cechlárová, K., T. Fleiner, D. F. Manlove, I. McBride, and E. Potpinková (2015). Modelling practical placement of trainee teachers to schools. *Central European Journal of Operations Research* 23(3), 547–562.
- Corcoran, S. P., J. L. Jennings, S. R. Cohodes, and C. Sattin-Bajaj (2018). Leveling the playing field for high school choice: Results from a field experiment of informational interventions. Technical report, National Bureau of Economic Research.
- Damm, A. P. (2014). Neighborhood quality and labor market outcomes: Evidence from quasi-random neighborhood assignment of immigrants. *Journal of Urban Economics* 79, 139–166.
- Delacrétaz, D. (2019, January). Stability in matching markets with sizes. Mimeo.
- Delacrétaz, D., S. D. Kominers, and A. Teytelboym (2016, November). Refugee resettlement. Mimeo.
- Dur, U. and M. U. Ünver (2015, March). Two-sided matching with balanced exchange: Tuition and worker exchanges.

- Echenique, F. and M. B. Yenmez (2007). A solution to matching with preferences over colleagues. *Games and Economic Behavior* 59(1), 46–71.
- Ehlers, L. (2008). Truncation strategies in matching markets. *Mathematics of Operations Research* 33(2), 327–335.
- Fernandez, M. (2017, October). Deferred acceptance and regret-free truth-telling: A characterization result. Technical report.
- Feywerda, J. and J. Gest (2016). Location, location: Refugee resettlement and integration outcomes in the United States. Mimeo.
- Fleiner, T. (2003). A fixed-point approach to stable matchings and some applications. *Mathematics of Operations Research* 28, 103–126.
- 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.
- Goto, M., N. Hashimoto, A. Iwasaki, Y. Kawasaki, S. Ueda, Y. Yasuda, and M. Yokoo (2014). Strategy-proof matching with regional minimum quotas. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*, pp. 1225–1232. International Foundation for Autonomous Agents and Multiagent Systems.
- Haake, C.-J. and B. Klaus (2009). Stability and Nash implementation in matching markets with couples. *Theory and Decision* 69(4), 537–554.
- Hafalir, I. E., M. B. Yenmez, and M. A. Yildirim (2013). Effective affirmative action in school choice. *Theoretical Economics* 8(2), 325–363.
- Hastings, J. S., R. Van Weelden, and J. Weinstein (2007). Preferences, information, and parental choice behavior in public school choice. Technical report, National Bureau of Economic Research.
- Hatfield, J. W. and P. Milgrom (2005). Matching with contracts. *American Economic Review* 95(4), 913–935.
- Jones, W. and A. Teytelboym (2017a). The International Refugee Match: A system that respects refugees’ preferences and the priorities of states. *Refugee Survey Quarterly* 36(2), 84–109.

- Jones, W. and A. Teytelboym (2017b). The Local Refugee Match: Aligning refugees' preferences with the capacities and priorities of localities. *Journal of Refugee Studies* 31(2), 152–178.
- Kamada, Y. and F. Kojima (2020, June). Fair matching under constraints: Theory and applications. Mimeo.
- Klaus, B. and F. Klijn (2005). Stable matchings and preferences of couples. *Journal of Economic Theory* 121(1), 75–106.
- Klaus, B., F. Klijn, and J. Massó (2007). Some things couples always wanted to know about stable matchings (but were afraid to ask). *Review of Economic Design* 11(3), 175–184.
- Kojima, F. (2012). School choice: Impossibilities for affirmative action. *Games and Economic Behavior* 75(2), 685–693.
- Kojima, F. (2015). Finding all stable matchings with couples. *Journal of Dynamics and Games* 2(3/4), 321–330.
- Kominers, S. D., A. Teytelboym, and V. P. Crawford (2017). An invitation to market design. *Oxford Review of Economic Policy* 33(4), 541–571.
- Martén, L., J. Hainmueller, and D. Hangartner (2019). Ethnic networks can foster the economic integration of refugees. *Proceedings of the National Academy of Sciences*, 201820345.
- McDermid, E. J. and D. F. Manlove (2010). Keeping partners together: algorithmic results for the hospitals/residents problem with couples. *Journal of Combinatorial Optimization* 19(3), 279–303.
- Mennle, T. and S. Seuken (2014). An axiomatic approach to characterizing and relaxing strategyproofness of one-sided matching mechanisms. In *Proceedings of the 15th ACM Conference on Economics and Computation (EC)*. Citeseer.
- Milgrom, P. and I. Segal (2020). Clock auctions and radio spectrum reallocation. *Journal of Political Economy* 128(1), 1–31.
- Moraga, J. F.-H. and H. Rapoport (2014). Tradable immigration quotas. *Journal of Public Economics* 115, 94–108.
- Nguyen, T., H. Nguyen, and A. Teytelboym (2019). Stability in matching markets with complex constraints. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pp. 61–61. ACM.

- 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.
- Pápai, S. (2000). Strategyproof assignment by hierarchical exchange. *Econometrica* 68(6), 1403–1433.
- Pápai, S. (2003). Strategyproof exchange of indivisible goods. *Journal of Mathematical Economics* 39(8), 931–959.
- Pápai, S. (2007). Exchange in a general market with indivisible goods. *Journal of Economic Theory* 132(1), 208–235.
- Pycia, M. and M. U. Ünver (2017). Incentive compatible allocation and exchange of discrete resources. *Theoretical Economics* 12(1), 287–329.
- Roth, A. E. (1984a). 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. (1984b). Stability and polarization of interests in job matching. *Econometrica* 52(1), 47–57.
- Roth, A. E. (2015, 9 March). Migrants aren't widgets. *Politico*. <https://www.politico.eu/article/migrants-arent-widgets-europe-eu-migrant-refugee-crisis/>.
- Roth, A. E. (2018). Marketplaces, markets, and market design. *American Economic Review* 108(7), 1609–58.
- Roth, A. E. and E. Peranson (1999). The redesign of the matching market for American physicians: Some engineering aspects of economic design. *American Economic Review* 89(4), 748–780.
- Schuck, P. H. (1997). Refugee burden-sharing: A modest proposal. *Yale Journal of International Law* 22, 243.
- Shapley, L. and H. Scarf (1974). On cores and indivisibility. *Journal of Mathematical Economics* 1(1), 23–37.
- Sönmez, T. (1997). Manipulation via capacities in two-sided matching markets. *Journal of Economic Theory* 77(1), 197–204.

- Sotomayor, M. (1996). A non-constructive elementary proof of the existence of stable marriages. *Games and Economic Behavior* 13(1), 135–137.
- Trapp, A., A. Teytelboym, A. Martinello, T. Andersson, and N. Ahani (forthcoming, 2020). Placement optimization in refugee resettlement. *Operations Research*.
- UNHCR (2019a, June). Global Resettlement Needs 2020. Technical report, United Nations High Commissioner for Refugees.
- UNHCR (2019b, June). Resettlement Fact Sheet January-June 2019. Technical report, United Nations High Commissioner for Refugees.
- UNHCR (2020). Global trends: Forced displacement in 2019. Technical report, United Nations High Commissioner for Refugees, Geneva.
- Westkamp, A. (2013). An analysis of the German university admissions system. *Economic Theory* 53(3), 561–589.
- Wiswall, M. and B. Zafar (2017). Preference for the workplace, investment in human capital, and gender. *The Quarterly Journal of Economics* 133(1), 457–507.
- Wu, Q. and A. E. Roth (2018). The lattice of envy-free matchings. *Games and Economic Behavior* 109, 201–211.

Appendix: Proofs

Proof of Proposition 1

Recall that μ^i is the current matching in Round i for the KTTC algorithm. We first verify that the output of the KTTC mechanism does not violate any capacity constraint; that is, for all $\ell \in L$, ℓ can accommodate $\mu^{KTTC}(\ell)$. Notice that, in any Round $i \geq 1$, a locality ℓ permanently rejects any family that ℓ cannot accommodate alongside $\mu^i(\ell)$. Therefore, if ℓ can accommodate $\mu^i(\ell)$, then ℓ can accommodate $\mu^{i+1}(\ell)$. Since ℓ can accommodate $\mu^1(\ell) = \emptyset$, it follows by induction that ℓ can accommodate the families currently matched to ℓ in every Round $i \geq 1$. As μ^{KTTC} is the current matching in the last round, ℓ can accommodate $\mu^{KTTC}(\ell)$. We next show that the KTTC mechanism is Pareto-efficient (PE) and strategy-proof (SP).

Proof of (PE) For every Round $i = 1, \dots, N$, let F^i be the set of families that have been permanently matched by the end of Round $i - 1$. (Note that $F^1 = \emptyset$ and $F^{N+1} = F$.) The proof proceeds by induction, with the following hypothesis: there does not exist any matching μ such that $\mu(f) \succeq_f \mu^i(f)$ for all $f \in F^i$ and $\mu(f) \succ_f \mu^i(f)$ for some $f \in F^i$. Our induction hypothesis trivially holds for $i = 1$ since $F^1 = \emptyset$. We now show that if our induction hypothesis holds for some $i = 1, \dots, N$, then it also holds for $i + 1$.

Towards a contradiction, suppose there exists a matching μ such that $\mu(f) \succeq_f \mu^{i+1}(f)$ for all $f \in F^{i+1}$ and $\mu(f) \succ_f \mu^{i+1}(f)$ for some $f \in F^{i+1}$. Note that, for all $f \in F^i$, $\mu^i(f) = \mu^{i+1}(f)$. If, for some $f \in F^i$, $\mu(f) \succ_f \mu^{i+1}(f)$, the induction hypothesis implies that there exists of a family $f' \in F^i$ such that $\mu^{i+1}(f') \succ_{f'} \mu(f')$, a contradiction. Therefore, $\mu^{i+1}(f) = \mu(f)$ for all $f \in F^i$. Note that, for all $f \in F^{i+1} \setminus F^i$, f is permanently matched to $\mu^{i+1}(f)$ in Round i ; therefore $\mu^{i+1}(f)$ is the locality f prefers among those that have not permanently rejected f . If, for some $f \in F^{i+1} \setminus F^i$, $\mu(f) \succ_f \mu^{i+1}(f)$, then $\mu(f)$ has permanently rejected f so it must be that $\mu(f)$ cannot accommodate f alongside $\mu^i(\mu(f))$. Therefore, there exists a family $f' \in F^i$ such that $\mu^i(f') \neq \mu(f')$, a contradiction.

By induction, there does not exist any matching μ such that $\mu(f) \succeq_f \mu^{N+1}(f)$ for all $f \in F^{N+1}$ and $\mu(f) \succ_f \mu^{N+1}(f)$ for some $f \in F^{N+1}$. As $\mu^{N+1} = \mu^{KTTC}$ and $F^{N+1} = F$, this implies that μ^{KTTC} is Pareto-efficient.

Proof of (SP) The result is implied by the fact that the KTTCE mechanism is strategy-proof (Theorem 3) and the fact that $\mu^{KTTCE} = \mu^{KTTC}$ when every family is endowed \emptyset (Proposition 3). \square

Proof of Theorem 1

The proof is by counterexample. There are four families, four localities, and one dimension. The endowment is

$$\mu^E = \begin{pmatrix} f_1 & f_2 & f_3 & f_4 \\ \ell_2 & \ell_2 & \ell_3 & \ell_4 \end{pmatrix}.$$

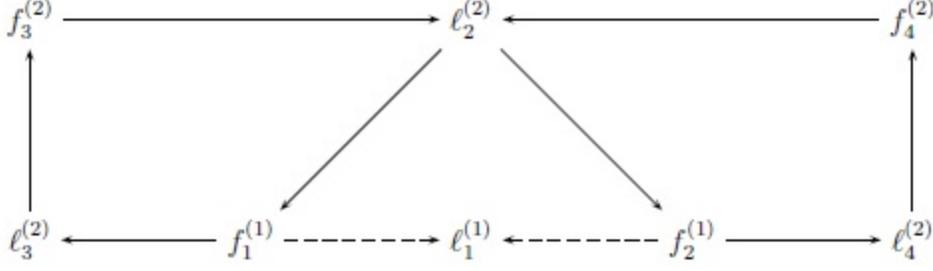


Figure 5: Counterexample for the proof of Theorem 1. $f \rightarrow l$: l is f 's first choice. $f \rightarrow l$: l is f 's second choice and f prefers l to its endowment. $l \rightarrow f$: l is f 's endowment. Superscripts denote sizes and capacities, respectively.

The preferences of families are as follows:

$$\succ_{f_1}: l_3, l_1, \mathbf{l_2}, \dots \quad \succ_{f_2}: l_4, l_1, \mathbf{l_2}, \dots \quad \succ_{f_3}: l_2, \mathbf{l_3}, \dots \quad \succ_{f_4}: l_2, \mathbf{l_4}, \dots,$$

where a family's endowment locality is denoted in boldface. The family sizes and locality capacities are

$$\nu =_{d_1} \begin{pmatrix} f_1 & f_2 & f_3 & f_4 \\ 1 & 1 & 2 & 2 \end{pmatrix} \quad \kappa =_{d_1} \begin{pmatrix} l_1 & l_2 & l_3 & l_4 \\ 1 & 2 & 2 & 2 \end{pmatrix}.$$

For ease of exposition, we illustrate this counterexample in Figure 5.

We show that there are two individually rational (IR) and chain-efficient (CE) matchings:

$$\mu = \begin{pmatrix} f_1 & f_2 & f_3 & f_4 \\ l_1 & l_4 & l_3 & l_2 \end{pmatrix} \quad \text{and} \quad \mu' = \begin{pmatrix} f_1 & f_2 & f_3 & f_4 \\ l_3 & l_1 & l_2 & l_4 \end{pmatrix}.$$

It is easy to check that matchings μ and μ' are IR and Pareto-efficient; hence they are IR and CE. It remains to show that no other IR matching is CE.

First, consider the case where f_1 is matched to l_3 . Then, f_3 must be matched to l_2 and f_4 must be matched to l_4 . Therefore, l_1 must be matched to f_2 , yielding μ' ; hence μ' is the only IR (and CE) matching where f_1 is matched to l_3 . Analogous reasoning allows us to conclude that μ is the only IR (and CE) matching where f_2 is matched to l_4 .

Second, consider the case where f_1 is matched to its endowment locality l_2 . Then, f_3 and f_4 must also be matched to their respective endowments, l_3 and l_4 . Hence, f_2 must be matched to either l_2 or l_1 . Matching f_2 to l_2 yields μ^E . Since f_2 prefers l_1 to l_2 and l_1 is not matched to any family, μ^E is wasteful, hence not CE. Matching f_2 to l_1 yields a matching that is not CE as it contains the Pareto-improving chain $f_1 \rightarrow l_3 \rightarrow f_3 \rightarrow l_2$. Therefore, there is no IR and CE matching where f_1 is matched to l_2 . Analogous reasoning allows us to reach the same conclusion for f_2 .

Third, consider the case where f_1 is matched to l_1 . Then, f_2 must be matched to either l_4 or l_2 , which by our previous argument yields either μ or μ^E . We therefore conclude that μ and μ' are the only two IR and CE matchings.

Suppose now that f_1 misreports its preferences by ranking l_1 below l_2 , i.e., $\succ'_{f_1}: l_3, \mathbf{l_2}, \dots$. Using analogous reasoning as for the true preference profile, it is easy to check that μ' is the

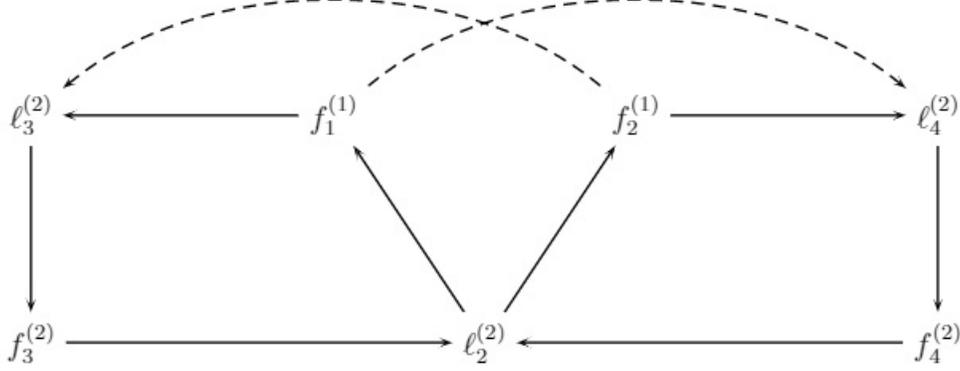


Figure 6: Counterexample for the proof of Proposition 2. $f \rightarrow l$: l is f 's first choice. $f \rightarrow l$: l is f 's second choice and f prefers l to its endowment. $l \rightarrow f$: l is f 's endowment. Superscripts denote sizes and capacities, respectively.

only IR and CE matching for the manipulated preference profile $(\succ'_{f_1}, \succ_{-f_1})$. Similarly, μ is the only IR and CE matching for the manipulated preference profile $(\succ'_{f_2}, \succ_{-f_2})$ where f_2 misreports its preferences by ranking l_1 below l_2 , i.e., $\succ'_{f_2}: l_4, \mathbf{l}_2, \dots$

We can now show that no IR and CE mechanism is strategy-proof. Let φ be an IR and CE mechanism. If all families report their preferences truthfully, then either $\varphi(\succ) = \mu$ or $\varphi(\succ) = \mu'$ because μ and μ' are the only two IR and CE matchings. If f_1 reports \succ'_{f_1} , then $\varphi(\succ'_{f_1}, \succ_{-f_1}) = \mu'$. Similarly, if f_2 reports \succ'_{f_2} , then $\varphi(\succ'_{f_2}, \succ_{-f_2}) = \mu$. If $\varphi(\succ) = \mu$, then $\varphi(\succ'_{f_1}, \succ_{-f_1})(f_1) \succ_{f_1} \varphi(\succ)(f_1)$ but if $\varphi(\succ) = \mu'$, then $\varphi(\succ'_{f_2}, \succ_{-f_2})(f_2) \succ_{f_2} \varphi(\succ)(f_2)$. Therefore φ is not strategy-proof. \square

Proof of Proposition 2

The proof is by counterexample. There are four families, three localities, and one dimension. The endowment is

$$\mu^E = \begin{pmatrix} f_1 & f_2 & f_3 & f_4 \\ l_2 & l_2 & l_3 & l_4 \end{pmatrix}.$$

The preferences of families are as follows:

$$\succ_{f_1}: l_3, l_4, \mathbf{l}_2, \dots \quad \succ_{f_2}: l_4, l_3, \mathbf{l}_2, \dots \quad \succ_{f_3}: l_2, \mathbf{l}_3, \dots \quad \succ_{f_4}: l_2, \mathbf{l}_4, \dots$$

where a family's endowment locality is denoted in boldface. The family sizes and locality capacities capacities are

$$\nu = d_1 \begin{pmatrix} f_1 & f_2 & f_3 & f_4 \\ 1 & 1 & 2 & 2 \end{pmatrix} \quad \kappa = d_1 \begin{pmatrix} l_2 & l_3 & l_4 \\ 2 & 2 & 2 \end{pmatrix}.$$

For ease of exposition, we illustrate this counterexample in Figure 6.

We show that there are exactly two matchings that Pareto dominate μ^E :

$$\mu = \begin{pmatrix} f_1 & f_2 & f_3 & f_4 \\ \ell_3 & \ell_3 & \ell_2 & \ell_4 \end{pmatrix} \quad \text{and} \quad \mu' = \begin{pmatrix} f_1 & f_2 & f_3 & f_4 \\ \ell_4 & \ell_4 & \ell_3 & \ell_2 \end{pmatrix}.$$

Matching μ (resp. μ') Pareto dominates μ^E since it matches family f_4 (resp. f_3) to its endowment and make the other three families better off. We now show that no other matching Pareto dominates μ^E . Towards a contradiction, suppose there exists a matching $\tilde{\mu}$ that Pareto dominates μ^E and is different from both μ and μ' . Consider the case where $\tilde{\mu}(f_3) \succ_{f_3} \mu^E(f_3)$. Then we must have that $\tilde{\mu}(f_3) = \ell_2$. Therefore, locality ℓ_2 cannot accommodate any other family so $\tilde{\mu}(f_4) = \ell_4$. In turn, locality ℓ_4 cannot accommodate any other family so $\tilde{\mu}(f_1) = \tilde{\mu}(f_2) = \ell_3$ and $\tilde{\mu} = \mu$. Analogously, we have that $\tilde{\mu}(f_4) \succ_{f_4} \mu^E(f_4)$ implies $\tilde{\mu} = \mu'$. Therefore, $\tilde{\mu}(f_3) = \ell_3$ and $\tilde{\mu}(f_4) = \ell_4$. Hence, neither ℓ_3 nor ℓ_4 can accommodate another family so $\tilde{\mu}(f_1) = \tilde{\mu}(f_2) = \ell_2$ and $\tilde{\mu} = \mu^E$, a contradiction.

We have established that μ and μ' are the only two matchings that Pareto dominate μ^E with respect to the true preferences. We now consider a preference manipulation by f_1 . Suppose that f_1 misreports its true preferences by ranking ℓ_4 below ℓ_2 , i.e., f_1 reports $\succ'_{f_1}: \ell_3, \ell_2, \dots$. We claim that if all other families report truthfully, then μ is the only matching that Pareto dominates μ^E with respect to the manipulated preference profile $(\succ'_{f_1}, \succ_{-f_1})$. Suppose again towards a contradiction that there exists a matching $\tilde{\mu}$ that Pareto dominates μ^E and is different from μ . There are three cases. First, if $\tilde{\mu}(f_3) \succ_{f_3} \mu^E(f_3)$, then $\tilde{\mu}(f_3) = \ell_2$ and $\tilde{\mu}(f_4) = \ell_4$. Therefore, we must have that $\tilde{\mu}(f_1) = \tilde{\mu}(f_2) = \ell_3$ so $\tilde{\mu} = \mu$, a contradiction. Second, if $\tilde{\mu}(f_4) \succ_{f_4} \mu^E(f_4)$, then $\tilde{\mu}(f_4) = \ell_2$ and $\tilde{\mu}(f_3) = \ell_3$. Since neither ℓ_2 nor ℓ_3 can accommodate another family, f_1 must be matched to a less-preferred locality than $\mu^E(f_1) = \ell_2$ (according to f_1 's manipulated report) hence $\tilde{\mu}$ does not Pareto dominate μ^E , a contradiction. Third, if $\tilde{\mu}(f_3) = \ell_3$ and $\tilde{\mu}(f_4) = \ell_4$, then $\tilde{\mu} = \mu^E$, a contradiction. We therefore conclude that there is no matching that Pareto dominates μ^E for the preference profile $(\succ'_{f_1}, \succ_{-f_1})$ and that is different from μ . By analogous reasoning, one can verify that if f_2 manipulates its preferences by reporting $\succ'_{f_2}: \ell_4, \ell_2, \dots$ and all other families report truthfully, then μ' is the only matching that Pareto dominates μ^E for the preference profile $(\succ'_{f_2}, \succ_{-f_2})$.

We can now show that there is no strategy-proof mechanism that Pareto improves upon μ^E . Let φ be a mechanism that Pareto improves upon μ^E . If all families report their preferences truthfully, then either $\varphi(\succ) = \mu$ or $\varphi(\succ) = \mu'$ because μ and μ' are the only two matchings that Pareto dominate μ^E for the true preference profile. If f_1 reports \succ'_{f_1} , then $\varphi(\succ'_{f_1}, \succ_{-f_1}) = \mu$. Similarly, if f_2 reports \succ'_{f_2} , then $\varphi(\succ'_{f_2}, \succ_{-f_2}) = \mu'$. If $\varphi(\succ) = \mu$, then $\varphi(\succ'_{f_2}, \succ_{-f_2})(f_2) \succ_{f_2} \varphi(\succ)(f_2)$ but if $\varphi(\succ) = \mu'$, then $\varphi(\succ'_{f_1}, \succ_{-f_1})(f_1) \succ_{f_1} \varphi(\succ)(f_1)$. Therefore, φ is not strategy-proof. \square

Proof of Theorem 2

The proof is by counterexample. There are four families, three localities, and two dimensions. The endowment is

$$\mu^E = \begin{pmatrix} f_1 & f_2 & f_3 & f_4 \\ \ell_2 & \ell_2 & \ell_3 & \ell_4 \end{pmatrix}.$$

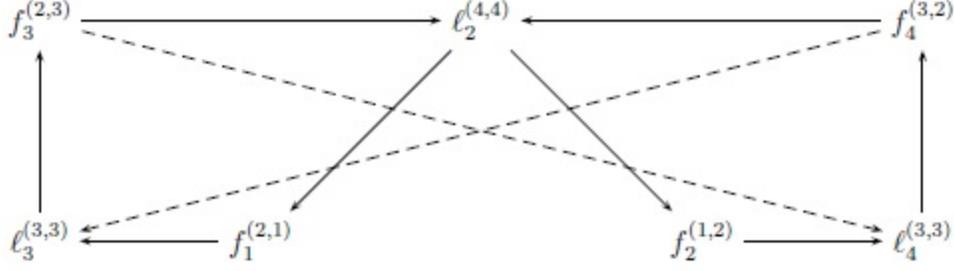


Figure 7: Counterexample for the proof of Theorem 2. $f \rightarrow l$: l is f 's first choice. $f \dashrightarrow l$: l is f 's second choice and f prefers l to its endowment. $l \rightarrow f$: l is f 's endowment. Superscripts denote sizes and capacities, respectively.

The preferences of families are as follows:

$$\succ_{f_1}: l_3, \mathbf{l_2}, \dots \quad \succ_{f_2}: l_4, \mathbf{l_2}, \dots \quad \succ_{f_3}: l_2, l_4, \mathbf{l_3}, \dots \quad \succ_{f_4}: l_2, l_3, \mathbf{l_4}, \dots$$

where a family's endowment locality is denoted in boldface. The family sizes and locality capacities are

$$\nu = \begin{matrix} & f_1 & f_2 & f_3 & f_4 \\ \begin{matrix} d_1 \\ d_2 \end{matrix} & \begin{pmatrix} 2 & 1 & 2 & 3 \\ 1 & 2 & 3 & 2 \end{pmatrix} \end{matrix} \quad \kappa = \begin{matrix} & l_2 & l_3 & l_4 \\ \begin{matrix} d_1 \\ d_2 \end{matrix} & \begin{pmatrix} 3 & 3 & 4 \\ 3 & 3 & 4 \end{pmatrix} \end{matrix}.$$

For ease of exposition, we illustrate this counterexample in Figure 7.

Observe that μ^E is non-wasteful but not chain-efficient as it contains three Pareto-improving chains:

$$(f_2, l_4, f_4, l_3, f_3, l_2), \quad (f_1, l_3, f_3, l_4, f_4, l_2), \quad \text{and} \quad (f_3, l_4, f_4, l_3).$$

Executing any one of these Pareto-improving chains on μ^E yields one of the following three matchings:

$$\mu = \begin{pmatrix} f_1 & f_2 & f_3 & f_4 \\ l_2 & l_4 & l_2 & l_3 \end{pmatrix}, \quad \mu' = \begin{pmatrix} f_1 & f_2 & f_3 & f_4 \\ l_3 & l_2 & l_4 & l_2 \end{pmatrix}, \quad \text{and} \quad \mu'' = \begin{pmatrix} f_1 & f_2 & f_3 & f_4 \\ l_2 & l_2 & l_4 & l_3 \end{pmatrix}.$$

We first show that μ , μ' , and μ'' are the only matchings that Pareto dominate the endowment μ^E .

Suppose towards a contradiction that there exists a matching $\tilde{\mu}$ that Pareto dominates μ^E . There are three cases.

First, consider the case where $\tilde{\mu}(f_3) = l_3$. Then, neither f_1 nor f_4 can be matched to l_3 ; hence $\tilde{\mu}(f_1) = l_2$ and $\tilde{\mu}(f_4) = l_4$. Therefore, we must have that $\tilde{\mu}(f_2) = l_2$. It follows that $\tilde{\mu} = \mu^E$, a contradiction.

Second, consider the case where $\tilde{\mu}(f_3) = l_2$. Then, f_2 cannot be matched to l_2 so $\tilde{\mu}(f_2) = l_4$. As a result, f_4 cannot be matched to either l_2 or l_4 , hence $\tilde{\mu}(f_4) = l_3$. It follows that f_1 cannot be matched to l_3 ; therefore $\tilde{\mu}(f_3) = l_2$ and $\tilde{\mu} = \mu$, a contradiction.

Third, consider the case where $\tilde{\mu}(f_3) = l_4$. Neither f_2 nor f_4 can be matched to l_4 ; therefore $\tilde{\mu}(f_2) = l_2$ and we must have that either $\tilde{\mu}(f_4) = l_2$ or $\tilde{\mu}(f_4) = l_3$. If $\tilde{\mu}(f_4) = l_2$,

then $\tilde{\mu}(f_1) = \ell_3$ and $\tilde{\mu} = \mu'$. If $\tilde{\mu}(f_4) = \ell_3$, then $\tilde{\mu}(f_1) = \ell_2$ and $\tilde{\mu} = \mu''$, a contradiction.

We have established that μ , μ' , and μ'' are the only matchings that Pareto dominate μ^E . Suppose that f_3 misreports its preferences by ranking ℓ_4 below ℓ_3 , i.e., by reporting $\succ'_{f_3}: \ell_2, \ell_3, \dots$, then μ is the unique matching that Pareto dominates μ^E for the manipulated preference profile $(\succ'_{f_3}, \succ_{-f_3})$. Similarly, suppose that f_4 misreports its preferences by ranking ℓ_3 below ℓ_4 , i.e., by reporting $\succ'_{f_4}: \ell_2, \ell_4, \dots$, then μ' is the unique matching that Pareto dominates μ^E for the manipulated preference profile $(\succ'_{f_4}, \succ_{-f_4})$.

We can now show that there is no strategy-proof mechanism that Pareto improves upon μ^E . Let φ be a mechanism that Pareto improves upon μ^E . If all families report their preferences truthfully, then $\varphi(\succ) \in \{\mu, \mu', \mu''\}$ because μ , μ' , and μ'' are the only three matchings that Pareto dominate μ^E for the true preference profile. If f_3 reports \succ'_{f_3} , then $\varphi(\succ'_{f_3}, \succ_{-f_3}) = \mu$. If f_4 reports \succ'_{f_4} , then $\varphi(\succ'_{f_4}, \succ_{-f_4}) = \mu'$. If $\varphi(\succ) \in \{\mu, \mu''\}$, then $\varphi(\succ'_{f_4}, \succ_{-f_4})(f_4) \succ_{f_4} \varphi(\succ)(f_4)$. If $\varphi(\succ) \in \{\mu', \mu''\}$, then $\varphi(\succ'_{f_3}, \succ_{-f_3})(f_3) \succ_{f_3} \varphi(\succ)(f_3)$. Therefore, φ is not strategy-proof. \square

Proof of Theorem 3

We first verify that the output of the KTTCE mechanism does not violate any capacity constraint; that is, for all $\ell \in L$, ℓ can accommodate $\mu^{KTTCE}(\ell)$. Note that, in any Round $i \geq 1$, the current matching μ^i is only updated when a feasible cycle is carried out. By definition, every family f that moves to locality ℓ as part of a feasible cycle can be accommodated by ℓ alongside all the families that remain at ℓ after the cycle is carried out; therefore, if ℓ can accommodate $\mu^i(\ell)$, then ℓ can accommodate $\mu^{i+1}(\ell)$. By assumption, the endowment does not violate any capacity constraint; hence, by induction, ℓ can accommodate the families currently matched to ℓ in every Round $i \geq 1$. As μ^{KTTCE} is obtained by carrying out feasible cycles from the current matching in the last round of the KTTCE algorithm, then ℓ can accommodate $\mu^{KTTCE}(\ell)$.

We next show that the KTTCE mechanism is individually rational (IR) and strategy-proof (SP).

Proof of (IR) Consider a family f , its endowment ℓ , and let i be the round of the KTTCE algorithm in which f is permanently matched. We need to show that $\mu^{KTTCE}(f) \succeq_f \ell$. By construction, any family that is not permanently matched is currently matched to its endowment; thus f is currently matched to ℓ at the start of Round i , i.e., $f \in \mu^i(\ell)$. Recall that in any round, locality ℓ can accommodate families currently matched to ℓ , in particular ℓ can accommodate $\mu^i(\ell)$. As a result, ℓ has not permanently rejected f . By definition, f points at its most preferred locality that has not permanently rejected f yet. Therefore, f points at a locality that f weakly prefers to ℓ . As f is permanently matched in Round i , f points at $\mu^{KTTCE}(f)$. Therefore, $\mu^{KTTCE}(f) \succeq_f \ell$, which is what IR requires.

Proof of (SP) Consider a family f with true preferences \succ_f , let \succ'_f be an alternative report, and fix the reports of all other families to \succ_{-f} . Denote by $\ell = \varphi^{KTTCE}(\succ_f, \succ_{-f})(f)$, respectively $\ell' = \varphi^{KTTCE}(\succ'_f, \succ_{-f})(f)$, the locality with which f is matched if it reports \succ_f , respectively \succ'_f . We need to show that $\ell \succeq_f \ell'$. Let $m = 1, \dots, N$, respectively $m' =$

$1, \dots, N$, be the round in which f gets permanently matched with report \succ_f , respectively \succ'_f .

Case 1: $m \leq m'$. Whether it reports \succ_f or \succ'_f , f is not permanently matched at the start of Round m . Let L_f^m be the set of localities that have not permanently rejected f at the start of Round m (this set is nonempty as we have already shown that a family is never permanently rejected by its endowment). Notice that, in any given round, f 's report does not impact whether or not a given locality permanently rejects f ; therefore L_f^m is the same whether f reports \succ_f or \succ'_f . In addition, permanent rejections are irreversible; therefore, with either report, f will be matched to a locality in L_f^m ; therefore, $\ell, \ell' \in L_f^m$. If f reports truthfully, it points at its most preferred locality in L_f^m and is permanently matched to it. Hence, ℓ is f 's most preferred locality in L_f^m , which implies that $\ell \succeq_f \ell'$.

Case 2: $m > m'$. As before, observe that, whether it reports \succ_f or \succ'_f , f is not permanently matched at the start of Round m' and let $L_f^{m'}$ be the set of localities that have not permanently rejected f at the start of Round m' . If f reports \succ'_f , f points at ℓ' in Round m' . A feasible cycle $f \rightarrow \ell' \rightarrow f_2 \rightarrow \ell_2 \rightarrow \dots \rightarrow f_n \rightarrow \ell_n \rightarrow f$ appears and f is permanently matched to ℓ' . By construction, in the KTTCE algorithm, the report of family f affects other families' pointing behavior only after f is permanently matched. Thus, as f is not matched until Round $m > m'$ when it reports \succ_f and every family and locality can be part of at most one cycle, none of f_2, \dots, f_n and $\ell', \ell_1, \dots, \ell_n$ are in a cycle in Round m' , hence none of them are permanently matched in Round m' .

We next show that none of f_2, \dots, f_n and $\ell', \ell_1, \dots, \ell_n$ are permanently matched before the start of Round m . By assumption, the cycle $f \rightarrow \ell' \rightarrow f_2 \rightarrow \ell_2 \rightarrow \dots \rightarrow f_n \rightarrow \ell_n \rightarrow f$ is feasible, which means that for all $j = 2, \dots, n-1$, ℓ_j can accommodate f_j alongside $\mu^{m'}(\ell_j) \setminus \{f_{j+1}\}$. Therefore, ℓ_j does not permanently reject f_j in any Rejection Stage before Round m . Analogously, ℓ_1 does not permanently reject f and ℓ_n does not permanently reject f_n in any Rejection Stage before Round m . Therefore, in Round m , all families f_2, \dots, f_n and all localities $\ell', \ell_2, \dots, \ell_n$ continue to point as they did in Round m' , meaning that f is permanently matched to ℓ' if f points at ℓ' . Since ℓ' has not permanently rejected f by Round m , f is permanently matched to either ℓ' or to a more preferred locality. Therefore, f can do no worse by reporting \succ_f than by reporting \succ'_f . \square

Proof of Proposition 3

We first show that all cycles that appear in the KTTCE algorithm are feasible. Towards a contradiction, suppose that an infeasible cycle $f_1 \rightarrow \ell_1 \rightarrow f_2 \rightarrow \ell_2 \rightarrow \dots \rightarrow f_n \rightarrow \ell_n \rightarrow f_1$ appears in some Round i . Then, there exists $j = 1, \dots, n$ such that ℓ_j cannot accommodate f_j alongside $\mu^i(\ell_j) \setminus \{f_{j+1}\}$ (letting $f_{n+1} = f_1$). As \emptyset can accommodate all families, $\ell_j \neq \emptyset$ so $\mu^E(\ell_j) = \emptyset$; hence f_{j+1} is not in ℓ_j 's endowment. Moreover, f_{j+1} is not permanently matched at the start of Round i as otherwise ℓ_j would not point at f_{j+1} ; therefore $f \in \mu^i(\emptyset)$, which implies $f \notin \mu^i(\ell_j)$. Then, $\mu^i(\ell_j) \setminus \{f_{j+1}\} = \mu^i(\ell_j)$; hence ℓ_j cannot accommodate f_j alongside $\mu^i(\ell_j)$. As $\mu^E(\ell_j) = \emptyset$, by construction all families in $\mu^i(\ell_j)$ have been permanently matched to ℓ_j before the start of Round i . Therefore, at the start of Round i , ℓ_j permanently rejects f_j as ℓ_j cannot accommodate f_j alongside all the families permanently matched to ℓ_j . We conclude that f_j does not point to ℓ_j in Round i , a contradiction.

We have showed that all cycles that appear in the KTTCE algorithm are feasible, which

implies that every round ends in the Matching Stage. Therefore, the KTTCE algorithm carries out all the cycles that appear in every round and the algorithm never enters the Rejection Stage. As a result, each round of the KTTCE algorithm coincides with the corresponding round of the KTTC algorithm. \square

Proof of Theorem 4

Consider an instance in which the sizes of families are monotonic, the priorities of localities are lexicographic, and the endowment is μ^E . Suppose that for this instance the KTTCE mechanism produces the endowment, i.e., $\mu^{KTTCE} = \mu^E$. Since the KTTCE mechanism does not Pareto improve upon this endowment, we need to show that μ^E is chain-efficient. Let N be the total number of rounds of the KTTCE algorithm. The fact that the KTTCE algorithm produces μ^E implies that the current matching is the same in every round: $\mu^E = \mu^1 = \mu^2 = \dots = \mu^N = \mu^{N+1} = \mu^{KTTCE}$. Let M^* be the set of matchings that can be obtained by starting from μ^E and carrying out exactly one Pareto-improving chain. Now μ^E is chain-efficient if and only if $M^* = \emptyset$; therefore it remains to show that $M^* = \emptyset$.

We proceed by induction with the following hypothesis: if a locality ℓ has permanently rejected a family f by the start of Round i , then for all $\mu^* \in M^*$, $\mu^*(f) \neq \ell$. Our inductive hypothesis trivially holds for $i = 1$ since no permanent rejection occurs before the start of Round 1. Assuming that the inductive hypothesis holds for some $i = 1, \dots, N$, we show that it holds for $i + 1$.

Consider a family f that has been permanently matched before the start of Round i . Recall that $\mu^E = \mu^i$, so f is permanently matched to its endowment $\mu^E(f)$; therefore f has been permanently rejected by all localities it prefers to $\mu^E(f)$. By the induction hypothesis, we have that, for all $\mu^* \in M^*$, $\mu^*(f) = \mu^E(f)$. We can therefore conclude that all families that are permanently matched before the start of Round i are matched to their endowment in all matchings contained in M^* .

Now consider a locality ℓ that permanently rejects a family f in Round i . We need to show that $\mu^*(f) \neq \ell$ for all $\mu^* \in M^*$. The fact that ℓ permanently rejects f in Round i implies that, at the start of Round i , f is not permanently matched and ℓ has not permanently rejected f yet. There are two cases in which ℓ can permanently reject f : (1) at the beginning of Round i , or (2) in the Rejection Stage of Round i .

Case 1: ℓ permanently rejects f at the beginning of Round i . By definition, ℓ cannot accommodate f alongside all the families permanently matched to ℓ . Since all families that are permanently matched before the start of Round i are matched to their endowment in all matchings contained in M^* , we have that all the families that are permanently matched to ℓ at the start of Round i are also matched to ℓ in all matchings contained in M^* . Therefore, $\mu^*(f) \neq \ell$ for all $\mu^* \in M^*$.

Case 2: ℓ permanently rejects f in the Rejection Stage of Round i . By definition, ℓ cannot accommodate f alongside $\mu^E(\ell) \setminus \{f'\}$ (where f' is the family at which ℓ is pointing, if any). Since a family is never permanently rejected by its endowment, we have that $\ell \neq \mu^E(f)$. By construction of the KTTCE algorithm, a locality does not point at any family if the locality has permanently rejected all families that are not permanently matched. Then the fact that ℓ has not permanently rejected f yet and that f is not permanently matched implies that ℓ does point at some family $f' \in F$. By construction, f' has the highest priority among all

families that are not permanently matched and have not been permanently rejected by ℓ . We consider two sub-cases: (2.1) $f' \notin \mu^E(\ell)$, and (2.2) $f \in \mu^E(\ell)$.

Sub-case 2.1: $f' \notin \mu^E(\ell)$. As priorities are lexicographic, all families in $\mu^E(\ell)$ have a higher priority than f' . The fact that ℓ points at f' implies that all families in $\mu^E(\ell)$ either have been permanently rejected by ℓ or have been permanently matched to ℓ . As a family cannot be permanently rejected by its endowment, all families in $\mu^E(\ell)$ are permanently matched to ℓ at the start of Round i . Recall that ℓ cannot accommodate f alongside $\mu^i(\ell) \setminus \{f'\}$. Since $\mu^i(\ell) = \mu^E(\ell)$ and $f' \notin \mu^E(\ell)$, ℓ cannot accommodate f alongside $\mu^i(\ell)$. Therefore, ℓ permanently rejects f at the beginning of Round i , contradicting our assumption that the permanent rejection occurs in the Rejection Stage.

Sub-case 2.2: $f' \in \mu^E(\ell)$. Towards a contradiction, suppose there exists a matching $\mu^* \in M^*$ such that $\mu^*(f) = \ell$. Recall that ℓ cannot accommodate f alongside $\mu^E(\ell) \setminus \{f'\}$; therefore ℓ cannot accommodate f alongside $\mu^E(\ell)$. This implies the existence of a family $\tilde{f} \in \mu^E(\ell)$ such that $\mu^*(\tilde{f}) \neq \ell$. In fact, because μ^* is obtained by carrying out exactly one Pareto-improving chain, there exists exactly one such family \tilde{f} . If $\tilde{f} \triangleright_\ell f'$, then the fact that ℓ points at f' implies that \tilde{f} is permanently matched to ℓ at the start of Round i . Following the argument in Sub-case 2.1, we must therefore have that $\mu^*(\tilde{f}) = \mu^E(\tilde{f}) = \ell$, a contradiction. Therefore, we have that $f' \triangleright_\ell \tilde{f}$ or $f' = \tilde{f}$. Since $f', \tilde{f} \in \mu^E(\ell)$, sizes are monotonic and priorities are lexicographic, by definition we have that $\nu_d^{f'} \geq \nu_d^{\tilde{f}}$ for all $d \in D$. Therefore, the fact that ℓ cannot accommodate f alongside $\mu^E(\ell) \setminus \{f'\}$ implies that ℓ cannot accommodate f alongside $\mu^E(\ell) \setminus \{\tilde{f}\}$, a contradiction. We conclude that $\mu^*(f) \neq \ell$ for all $\mu^* \in M^*$.

We can now conclude that our induction hypothesis holds at the beginning of Round $i+1$. By induction, if a locality ℓ permanently rejects a family f at some point in the KTTCE algorithm, then, for all $\mu^* \in M^*$, $\mu^*(f) \neq \ell$. Therefore, every family matched to its endowment at the end of the algorithm is matched to its endowment under any matching $\mu^* \in M^*$. By assumption, all families are matched to their respective endowments at the end of the KTTCE algorithm, meaning that M^* does not contain any matching other than μ^E . As $\mu^E \notin M^*$ by definition, we conclude that $M^* = \emptyset$ as desired. \square

Proof of Proposition 4

Suppose that sizes are monotonic and the priority profile is aligned. We introduce the Sequential Deferred Acceptance (SDA) algorithm (Algorithm 6) and show that it produces a stable matching. We show that the alignment of the priority profile allows us to divide the families into groups such that (i) all families in the same group have the same size and (ii) any two groups can be compared in terms of priority in the sense that all families in one group have a higher priority at all localities than all families in the other group. The SDA algorithm considers one of these groups at a time in order of priority and runs the family-proposing Deferred Acceptance (DA) algorithm for the families in that group, considering the capacities that remain after families in higher-priority groups have been permanently matched.

Claim 1. *The Sequential Deferred Acceptance algorithm produces a stable matching μ^{SDA} .*

First, if there is a directed cycle (f_1, f_2, \dots, f_n) in graph \mathbb{G}^1 , then, for every $\ell \in L \setminus \{\emptyset\}$, $f_1 \triangleright_\ell f_2 \triangleright_\ell \dots \triangleright_\ell f_n \triangleright_\ell f_1$, a contradiction. Therefore, \mathbb{G}^1 is an acyclic directed graph. By

Algorithm 6: SEQUENTIAL DEFERRED ACCEPTANCE

Construct a directed graph \mathbb{G}^1 as follows. Each of the $|F|$ vertices represents a family. For every pair of families (f, f') , let there be a directed edge from f to f' if $\nu^f \neq \nu^{f'}$ and $f \triangleright_\ell f'$ for all $\ell \in L$.

For every locality ℓ , set a *counter* $c_\ell^1 = \kappa^\ell$.

Round $i \geq 1$

Let \tilde{F}^i be the set of families at which no other family is pointing in graph \mathbb{G}^i .

Permanently match the families in \tilde{F}^i to the localities using the family-proposing Deferred Acceptance algorithm setting the capacity of every locality ℓ to its counter c_ℓ^i .

If all families have been permanently matched, end.

Otherwise, construct \mathbb{G}^{i+1} by removing from \mathbb{G}^i all vertices representing families in \tilde{F}^i and all edges adjacent to them. For every $\ell \in L$, let \tilde{F}_ℓ^i be the set of families that have been permanently matched to ℓ in Round i . Update the counter of every locality ℓ as follows: $c_\ell^{i+1} = c_\ell^i - \sum_{f \in \tilde{F}_\ell^i} \nu^f$. Continue to Round $i + 1$.

construction, for all $i > 1$, \mathbb{G}^i is an directed acyclic graph for since \mathbb{G}^i is constructed from \mathbb{G}^{i-1} by removing some vertices (families) and edges. Therefore, in every Round i , the set of families \tilde{F}^i at which no family is pointing is nonempty. Therefore, at least one family gets permanently matched in every round; hence the algorithm ends after at most $|F|$ rounds.

Second, suppose that in Round i , there exist $f, f' \in \tilde{F}^i$ such that $\nu^f \neq \nu^{f'}$. Since the priority profile is aligned, either $f \triangleright_\ell f'$ for all $\ell \in L \setminus \{\emptyset\}$ or $f' \triangleright_\ell f$ for all $\ell \in L \setminus \{\emptyset\}$. Therefore, either f points at f' , in which case $f' \notin \tilde{F}^i$, or f' points at f , in which case $f \notin \tilde{F}^i$, a contradiction. We conclude that in every Round i we have that $\nu^f = \nu^{f'}$ for all families $f, f' \in \tilde{F}^i$.

We now show that the matching μ^{SDA} produced by the SDA algorithm is stable. Consider any family f and any locality ℓ such that $\ell \succ_f \mu^{SDA}(f)$. We need to show that ℓ cannot accommodate f alongside $\hat{F}_\ell^f \cap \mu^{SDA}(\ell)$.

Suppose that family f is permanently matched to $\mu^{SDA}(f)$ in Round i , i.e., $f \in \tilde{F}^i$. Consider the Round i “submarket”, in which the DA algorithm permanently matches the families in \tilde{F}^i to the localities and the counter of each locality ℓ is c_ℓ^i . By construction, all families in the Round i submarket have the same size, i.e., $\nu^{f'} = \nu^f$ for all $f' \in \tilde{F}^i$. Therefore, for each locality, there is a maximum number of families that the locality can accommodate, which makes the submarket isomorphic to a school choice problem. As the DA algorithm produces a stable matching in the school choice problem (Abdulkadiroğlu and Sönmez, 2003), the DA algorithm produces a stable matching in the Round i submarket. Therefore, f and ℓ do not form a blocking pair in the Round i submarket. As $\ell \succ_f \mu^{SDA}(f)$ by assumption, it follows that ℓ 's Round i counter does not allow ℓ to accommodate f alongside higher-priority families permanently matched to ℓ in Round i , i.e., alongside $\hat{F}_\ell^f \cap \tilde{F}_\ell^i$. By construction, all

families that have been permanently matched to ℓ before Round i , i.e., all families in $\cup_{j=1}^{i-1} \widetilde{F}^j$, have a higher priority at ℓ than f . Hence, ℓ cannot accommodate f alongside all families that have a higher priority than f and with which ℓ has been permanently matched by the end of Round i , i.e., ℓ cannot accommodate f alongside $\widehat{F}_\ell^f \cap (\cup_{j=1}^i \widetilde{F}_\ell^j)$. By construction, all families that are permanently matched to ℓ in any round of the SDA algorithm are matched to ℓ at μ^{SDA} , which implies that $\cup_{j=1}^i \widetilde{F}_\ell^j \subseteq \mu^{SDA}(\ell)$. We therefore conclude that ℓ cannot accommodate f alongside $\widehat{F}_\ell^f \cap \mu^{SDA}(\ell)$, as required. \square

Proof of Proposition 5

Consider the following modification of the KDA algorithm. In each round, every locality ℓ tentatively accepts a proposing family f if ℓ can *accommodate* (instead of *weakly accommodate*) f alongside all families with a higher priority than f at ℓ that are proposing to ℓ or have been permanently rejected by ℓ . To see that this modified KDA algorithm produces the family-optimal envy-free matching, follow the proof of Theorem 5 verbatim having replaced “weakly accommodate” with “accommodate”. \square

Proof of Proposition 6

Let μ be a matching that is non-wasteful and not stable. We need to show that μ is not weakly envy-free. As μ is not stable, there exists a family f and a locality ℓ such that $\ell \succ_f \mu(f)$ and ℓ can accommodate f alongside $\widehat{F}_\ell^f \cap \mu(\ell)$, that is, for all $d \in D$,

$$\nu_d^f + \sum_{g \in (\widehat{F}_\ell^f \cap \mu(\ell))} \nu_d^g \leq \kappa_d^\ell. \quad (1)$$

However, as μ is non-wasteful, there exists $d \in D$ such that

$$\nu_d^f + \sum_{g \in \mu(\ell)} \nu_d^g > \kappa_d^\ell. \quad (2)$$

Consider now the families in $\mu(\ell) \setminus \widehat{F}_\ell^f$ whose size in dimension d is at least one. Equations (1) and (2) imply that the set of such families is nonempty. Let f' be the lowest-priority family in $\mu(\ell) \setminus \widehat{F}_\ell^f$ whose size in dimension d is at least one (equivalently, f' is the lowest-priority family in $\mu(\ell)$ such that $\nu_d^{f'} > 0$). Then,

$$\nu_d^f + \nu_d^{f'} + \sum_{g \in (\widehat{F}_\ell^{f'} \cap \mu(\ell))} \nu_d^g = \nu_d^f + \sum_{g \in \mu(\ell)} \nu_d^g > \kappa_d^\ell. \quad (3)$$

Inequality (3) and the fact that $\nu_d^{f'} > 0$ imply that ℓ cannot weakly accommodate f' alongside $\{f\} \cup (\widehat{F}_\ell^{f'} \cap \mu(\ell))$. Note that all families in $\{f\} \cup (\widehat{F}_\ell^{f'} \cap \mu(\ell))$ have a higher priority at ℓ than f' and weakly prefer ℓ to the locality to which they are matched, i.e.,

$$(\{f\} \cup (\widehat{F}_\ell^{f'} \cap \mu(\ell))) \subseteq \{g \in F : g \triangleright_\ell f' \text{ and } \ell \succeq_g \mu(g)\},$$

so ℓ cannot weakly accommodate f' alongside $\{g \in F : g \succ_{\ell} f' \text{ and } \ell \succeq_g \mu(g)\}$. Since $\ell \succ_f \mu(f)$ and $f \succ_{\ell} f'$, f strongly envies f' and μ is not weakly envy-free. \square

Proof of Theorem 5

We first verify that the output of the KDA algorithm does not violate any capacity constraint; that is, we show that, for all $\ell \in L$, ℓ can accommodate $\mu^{KDA}(\ell)$. Consider a locality $\ell \in L$. By construction, all families in $\mu^{KDA}(\ell)$ propose to and are tentatively accepted by ℓ in the last round of the KDA algorithm. Therefore, ℓ can weakly accommodate every family $f \in \mu^{KDA}(\ell)$ alongside $\mu^{KDA}(\ell) \cap \widehat{F}_{\ell}^f$. Towards a contradiction, suppose that ℓ cannot accommodate $\mu^{KDA}(\ell)$. Then, there exists $d \in D$ such that $\sum_{f \in \mu^{KDA}(\ell)} \nu_d^f > \kappa_d^{\ell}$. Let $g \in F$ be the lowest-priority family such that $g \in \mu^{KDA}(\ell)$ and $\nu_d^g > 0$. (Such a family exists since $\sum_{f \in \mu^{KDA}(\ell)} \nu_d^f > \kappa_d^{\ell} \geq 0$.) Then,

$$\nu_d^g + \sum_{f \in \mu^{KDA}(\ell) \cap \widehat{F}_{\ell}^g} \nu_d^f > \kappa_d^{\ell} \quad \text{and} \quad \nu_d^g > 0;$$

therefore, ℓ cannot weakly accommodate g alongside $\mu^{KDA}(\ell) \cap \widehat{F}_{\ell}^g$ (see Definition 9), a contradiction.

We next show that μ^{KDA} is weakly envy-free; towards a contradiction, suppose the contrary. Then, by definition, there exists a family f' matched to a locality ℓ' ($\mu^{KDA}(f') = \ell'$) such that ℓ' cannot weakly accommodate f' alongside $\widehat{G} = \{g \in G : g \succ_{\ell'} f' \text{ and } \ell' \succeq_g \mu^{KDA}(g)\}$. Suppose that the KDA algorithm terminates in Round N . Let $\widehat{R}_{f'}^N(\ell')$ denote the set of families that have a higher priority than f' at ℓ' and propose in Round N to either ℓ' or a less-preferred locality (because they have already been permanently rejected by ℓ' in a previous round). Since all families are matched to the locality to which they propose in Round N , $\widehat{R}_{f'}^N(\ell')$ is also the set of families that have a higher priority than f' at ℓ' and are matched to either ℓ' or to less-preferred localities at the end of the algorithm; therefore $\widehat{R}_{f'}^N(\ell') = \widehat{G}$. Since ℓ' tentatively accepts f' in Round N , ℓ' can weakly accommodate f' alongside $\widehat{R}_{f'}^N(\ell') = \widehat{G}$, a contradiction.

We finally show that μ^{KDA} dominates all other weakly envy-free matchings. Towards a contradiction, suppose that there exists a weakly envy-free matching μ such that $\mu(f_1) \succ_{f_1} \mu^{KDA}(f_1)$ for some $f_1 \in F$. We proceed by induction. Our assumption implies that $\mu(f_1)$ permanently rejects f_1 in some Round i_1 of the KDA algorithm. For the induction argument, suppose that for some $n \in \mathbb{Z}_{>0}$, there exists a family f_n such that $\mu(f_n)$ permanently rejects f_n in some Round i_n of the KDA algorithm. We want to show that there exists a family f_{n+1} such that $\mu(f_{n+1})$ permanently rejects f_{n+1} in some Round $i_{n+1} < i_n$. By construction, $\mu(f_n)$ cannot weakly accommodate f_n alongside $\widehat{R}_{f_n}^{i_n}(\mu(f_n))$ and all families in $\widehat{R}_{f_n}^{i_n}(\mu(f_n))$ are either matched to $\mu(f_n)$ or a less-preferred locality at μ^{KDA} . Suppose that, at μ , all families in $\widehat{R}_{f_n}^{i_n}(\mu(f_n))$ are either matched to $\mu(f_n)$ or a less-preferred locality. Then, μ cannot be weakly envy-free since $\mu(f_n)$ cannot weakly accommodate f_n alongside $\widehat{R}_{f_n}^{i_n}(\mu(f_n))$.⁴¹ Therefore, there exists a family $f_{n+1} \in \widehat{R}_{f_n}^{i_n}(\mu(f_n))$ such that $\mu(f_{n+1}) \succ_{f_{n+1}} \mu(f_n)$. Since f_{n+1} proposes

⁴¹Recall that part (iii) of Definition 10 implies parts (i) and (ii) of Definition 10.

to $\mu(f_n)$ in some Round $j \leq i_n$ of the KDA algorithm, $\mu(f_{n+1})$ permanently rejects f_{n+1} in some Round $i_{n+1} < i_n$. Iterating this argument inductively, we find that some family f_m is permanently rejected by $\mu(f_m)$ in Round i_m such that $i_m < 1$, which is impossible. \square

Proof of Theorem 6

We first verify that the output of the TKDA mechanism does not violate any capacity constraint; that is, for all $\ell \in L$, ℓ can accommodate $\mu^{TKDA}(\ell)$. Consider a locality $\ell \in L$. By construction, all families in $\mu^{TKDA}(\ell)$ propose to and are tentatively accepted by ℓ in the last round of the TKDA algorithm. Hence, for any $f \in \mu^{TKDA}(\ell)$,

$$\theta_\ell^f \geq |\mu^{TKDA}(\ell) \cap \widehat{F}_\ell^f| + 1 \geq 1.$$

If ℓ cannot weakly accommodate f alongside $\mu^{TKDA}(\ell) \cap \widehat{F}_\ell^f$, then by construction (see Algorithm 5) $\theta_\ell^f = 0$, which contradicts the fact that $\theta_\ell^f \geq 1$. Therefore, ℓ can weakly accommodate every family $f \in \mu^{TKDA}(\ell)$ alongside $\mu^{TKDA}(\ell) \cap \widehat{F}_\ell^f$. As we showed in the proof of Theorem 5, this implies that ℓ can accommodate $\mu^{TKDA}(\ell)$.

We now prove that the TKDA mechanism is strategy-proof and weakly envy-free.

TKDA is strategy-proof

Consider a locality $\ell \in L$ and a subset of families $G \subseteq F$. To simplify notation, let us define, for every $f \in F$, $\widehat{G}_\ell^f = G \cap \widehat{F}_\ell^f$ to be the families in G that have a higher priority for ℓ than f . We also denote by $\theta_\ell^f(G)$ the threshold of family f at locality ℓ if, in some round of the TKDA algorithm, families in G propose to ℓ . (That is, $\theta_\ell^f(G)$ is obtained by running the Threshold Calculator defined in Algorithm 5 with $\Pi_\ell = G$.)

We define the *choice function* of locality ℓ , $C_\ell : 2^F \rightarrow 2^F$, as follows: for every $G \subseteq F$, $C_\ell(G) = \{f \in G : |\widehat{G}_\ell^f| + 1 \leq \theta_\ell^f(G)\}$. That is, $C_\ell(G)$ contains the families that ℓ does *not* permanently reject if families in G propose to ℓ in some round of the TKDA algorithm. One way to interpret the choice function is that locality ℓ “chooses” the families in $C_\ell(G)$ when it receives proposals from all families in G .

We now fix a locality $\ell \in L$ and two subsets of families $G \subseteq H \subseteq F$ and define two properties of the choice function:

- *Substitutability (S)*: $C_\ell(H) \cap G \subseteq C_\ell(G)$, and
- *Cardinal Monotonicity (CM)*: $|C_\ell(G)| \leq |C_\ell(H)|$.

Hatfield and Milgrom (2005) analyze properties of the Deferred Acceptance (DA) algorithm, in which localities (“hospitals” in their terminology) have choice functions. In each round of the DA algorithm, families propose to their most preferred locality that has not permanently rejected them yet. Localities tentatively accept or permanently reject proposals based on their choice functions, i.e., if locality ℓ receives proposals from families in G , families in $C_\ell(G)$ are tentatively accepted and families in $G \setminus C_\ell(G)$ are permanently rejected.

Theorems 3 and 11 of Hatfield and Milgrom (2005) imply that the DA algorithm is strategy-proof for families if the choice function of all localities satisfy the (S) and (CM)

conditions. By construction, the TKDA algorithm in our setting corresponds to deferred acceptance in the [Hatfield and Milgrom \(2005\)](#) setting with the choice function we have just defined. Therefore, in order to show that the TKDA algorithm is strategy-proof, it remains to prove that the choice function of every locality satisfies the (S) and (CM) conditions.

The following lemma shows three properties of the choice function, which we subsequently use to prove that the choice function satisfies the (S) and (CM) conditions.

Lemma 1. *For every $G \subseteq H \subseteq F$, every $f, g \in F$, and every $\ell \in L$:*

- (i) $\theta_\ell^f(G) = \infty$ if and only if $\theta_\ell^f(H) = \infty$,
- (ii) If $\theta_\ell^f(H) \in \mathbb{Z}_{>0}$, then $\theta_\ell^f(G) \leq \theta_\ell^f(H) \leq \theta_\ell^f(G) + |\widehat{H}_\ell^f| - |\widehat{G}_\ell^f|$, and
- (iii) If $g \triangleright_\ell f$ and $\theta_\ell^f(G) \neq \infty$, then
 - $\theta_\ell^g(G) \geq \theta_\ell^f(G)$, and
 - $f \in C_\ell(G)$ implies that $g \in G \Leftrightarrow g \in C_\ell(G)$.

Part (i) of Lemma 1 states that if the threshold of family f at ℓ is infinite, then it remains so no matter which families are proposing to ℓ . Part (ii) states that if the threshold of family f at ℓ is non-zero and finite, then removing some families from the set of families proposing to ℓ may reduce f 's threshold at ℓ by at most the number of proposing families that were removed. Part (iii) states that if the threshold of family f at ℓ is finite, then any family g with a higher priority than f at ℓ has a weakly larger threshold than f and is chosen by ℓ whenever f is chosen by ℓ .

Proof of (S) Consider any family $f \in C_\ell(H) \cap G$; we need to show that $f \in C_\ell(G)$. The fact that $f \in C_\ell(H)$ implies that $\theta_\ell^f(H) \neq 0$ and, if $\theta_\ell^f(H) = \infty$, Lemma 1(i) implies that $\theta_\ell^f(G) = \infty$, hence $f \in C_\ell(G)$. It remains to show that $f \in C_\ell(H)$ whenever $\theta_\ell^f(H) \in \mathbb{Z}_{>0}$. In that case, using Lemma 1(ii), we have that

$$\theta_\ell^f(H) \leq \theta_\ell^f(G) + |\widehat{H}_\ell^f| - |\widehat{G}_\ell^f|.$$

Since $f \in C_\ell(G)$, the definition of a choice function implies that

$$\theta_\ell^f(H) \geq |\widehat{H}_\ell^f| + 1.$$

Combining the two inequalities yields $\theta_\ell^f(G) \geq |\widehat{G}_\ell^f| + 1$ so $f \in C_\ell(G)$, as required.

Proof of (CM) We need to show that $|C_\ell(G)| \leq |C_\ell(H)|$. Let $m = |H \setminus G|$ and arbitrarily label the families in $H \setminus G$ such that $H = G \cup \{f_1, \dots, f_m\}$.

Claim 2. *For every $i = 1, \dots, m$, $|C_\ell(G \cup \{f_1, \dots, f_{i-1}\})| \leq |C_\ell(G \cup \{f_1, \dots, f_{i-1}, f_i\})|$.*

Claim 2 implies that

$$|C_\ell(G)| \leq |C_\ell(G \cup \{f_1\})| \leq |C_\ell(G \cup \{f_1, f_2\})| \leq \dots \leq |C_\ell(G \cup \{f_1, \dots, f_m\})| = |C_\ell(H)|,$$

which implies the desired result. Therefore, it remains to prove Claim 2.

Fix some $i = 1, \dots, m$ and define $G' = G \cup \{f_1, \dots, f_{i-1}\}$ and $H' = \{f_1, \dots, f_{i-1}, f_i\}$. We need to show that $|C_\ell(G')| \leq |C_\ell(H')|$. We have that

$$|C_\ell(H')| - |C_\ell(G')| = |C_\ell(H') \setminus C_\ell(G')| - |C_\ell(G') \setminus C_\ell(H')|,$$

so we need to show that

$$|C_\ell(H') \setminus C_\ell(G')| \geq |C_\ell(G') \setminus C_\ell(H')|. \quad (4)$$

By definition, $H' = G' \cup \{f_i\}$ and, by (S), $C_\ell(H') \cap G' \subseteq C_\ell(G')$; therefore,

$$C_\ell(H') \setminus C_\ell(G') \subseteq \{f_i\}.$$

There are two cases: $f_i \notin C_\ell(H')$ and $f_i \in C_\ell(H')$.

Case 1: $f_i \notin C_\ell(H')$. In this case, $C_\ell(H') \setminus C_\ell(G') = \emptyset$ so, by inequality (4), we need to show that $C_\ell(G') \setminus C_\ell(H') = \emptyset$. Towards a contradiction, suppose to the contrary that there exists a family $g \in C_\ell(G') \setminus C_\ell(H')$. Since $g \in C_\ell(G')$, by the definition of the choice function, we have that

$$\theta_\ell^g(G') \geq |\widehat{G}'^g| + 1, \quad (5)$$

but, since $g \notin C_\ell(H')$, we also have that

$$\theta_\ell^g(H') < |\widehat{H}'^g| + 1. \quad (6)$$

We now consider two subcases: $g \triangleright_\ell f_i$ and $f_i \triangleright_\ell g$.

Sub-case 1.1: $g \triangleright_\ell f_i$. In this case, as $H' = G' \cup \{f_i\}$, $\widehat{G}'^g = \widehat{H}'^g$ so $|\widehat{G}'^g| = |\widehat{H}'^g|$ and, as a family's threshold only depends on higher-priority families (Algorithm 5), $\theta_\ell^g(G') = \theta_\ell^g(H')$. Combining these observations with inequalities (5) and (6) yields

$$\theta_\ell^g(G') \geq |\widehat{G}'^g| + 1 = |\widehat{H}'^g| + 1 > \theta_\ell^g(H') = \theta_\ell^g(G'),$$

a contradiction.

Sub-case 1.2: $f_i \triangleright_\ell g$. On the one hand, inequality (6) implies that $\theta_\ell^g(H') \neq \infty$ so we can apply Lemma 1(i) to obtain that $\theta_\ell^g(G') \neq \infty$ and Lemma 1(iii) to obtain that

$$\theta_\ell^{f_i}(G') \geq \theta_\ell^g(G'). \quad (7)$$

On the other hand, by the assumption that $f_i \notin C_\ell(H')$ in Case 1, we have that

$$\theta_\ell^{f_i}(H') < |\widehat{H}'^{f_i}| + 1. \quad (8)$$

As $H' = G' \cup \{f_i\}$, $\widehat{G}'^{f_i} = \widehat{H}'^{f_i}$ so $|\widehat{G}'^{f_i}| = |\widehat{H}'^{f_i}|$ and, as a family's threshold only depends on higher-priority families (Algorithm 5), $\theta_\ell^{f_i}(G') = \theta_\ell^{f_i}(H')$. Combining these observations with inequality (8) yields

$$\theta_\ell^{f_i}(G') = \theta_\ell^{f_i}(H') < |\widehat{H}'^{f_i}| + 1 = |\widehat{G}'^{f_i}| + 1. \quad (9)$$

Combining inequality (9) with the fact that $|\widehat{G}_\ell^{f_i}| \leq |\widehat{G}_\ell^{g_i}|$ (as $f_i \triangleright g_i$) and inequality (5) yields

$$\theta_\ell^{f_i}(G') < |\widehat{G}_\ell^{f_i}| + 1 \leq |\widehat{G}_\ell^{g_i}| + 1 \leq \theta_\ell^{g_i}(G').$$

We conclude that $\theta_\ell^{f_i}(G') < \theta_\ell^{g_i}(G')$, which contradicts inequality (7).

Case 2: $f_i \in C_\ell(H')$. In this case, we have that $C_\ell(H') \setminus C_\ell(G') = \{f_i\}$; hence, by inequality (4), we need to show that $|C_\ell(G') \setminus C_\ell(H')| \leq 1$. Towards a contradiction, suppose that there exist two distinct families $g_1, g_2 \in C_\ell(G') \setminus C_\ell(H')$. Without loss of generality, we assume that $g_1 \triangleright_\ell g_2$. Since $g_1, g_2 \in C_\ell(G')$, by the definition of the choice function, we have that

$$\theta_\ell^{g_1}(G') \geq |\widehat{G}_\ell^{g_1}| + 1 \quad \text{and} \quad \theta_\ell^{g_2}(G') \geq |\widehat{G}_\ell^{g_2}| + 1, \quad (10)$$

but, since $g_1, g_2 \notin C_\ell(H')$, we also have that

$$\theta_\ell^{g_1}(H') < |\widehat{H}_\ell^{g_1}| + 1 \quad \text{and} \quad \theta_\ell^{g_2}(H') < |\widehat{H}_\ell^{g_2}| + 1. \quad (11)$$

We now consider two subcases: $g_1 \triangleright_\ell f_i$ and $f_i \triangleright_\ell g_1$.

Sub-case 2.1: $g_1 \triangleright_\ell f_i$. In this case, as $H' = G' \cup \{f_i\}$, $\widehat{G}_\ell^{g_1} = \widehat{H}_\ell^{g_1}$ so $|\widehat{G}_\ell^{g_1}| = |\widehat{H}_\ell^{g_1}|$ and, as a family's threshold only depends on higher-priority families (Algorithm 5), $\theta_\ell^{g_1}(G') = \theta_\ell^{g_1}(H')$. Combining these observations with inequalities (10) and (11) yields

$$\theta_\ell^{g_1}(G') \geq |\widehat{G}_\ell^{g_1}| + 1 = |\widehat{H}_\ell^{g_1}| + 1 > \theta_\ell^{g_1}(H') = \theta_\ell^{g_1}(G'),$$

a contradiction.

Sub-case 2.2: $f_i \triangleright_\ell g_1$. Inequality (11) implies that $\theta_\ell^{g_2}(H') \neq \infty$ so we can apply Lemma 1(i) to obtain that $\theta_\ell^{g_2}(G') \neq \infty$ and Lemma 1(iii) to obtain that $\theta_\ell^{g_1}(G') \geq \theta_\ell^{g_2}(G')$. Moreover, as $g_1 \triangleright_\ell g_2$ and $g_1 \in G$, we have that $\widehat{G}_\ell^{g_1} \subset \widehat{G}_\ell^{g_2}$ so $|\widehat{G}_\ell^{g_1}| < |\widehat{G}_\ell^{g_2}|$. Combining these observations with inequality (10) yields

$$\theta_\ell^{g_1}(G') \geq \theta_\ell^{g_2}(G') \geq |\widehat{G}_\ell^{g_2}| + 1 > |\widehat{G}_\ell^{g_1}| + 1.$$

We therefore conclude that

$$\theta_\ell^{g_1}(G') \geq |\widehat{G}_\ell^{g_1}| + 2. \quad (12)$$

Since $\theta_\ell^{g_1}(H') \neq \infty$, we have two cases to consider: $\theta_\ell^{g_1}(H') = 0$ and $\theta_\ell^{g_1}(H') \in \mathbb{Z}_{>0}$.

Sub-sub-case 2.2.1: $\theta_\ell^{g_1}(H') = 0$. In this case, by the definition of thresholds (Algorithm 5), there exists a family $h \in F$ such that (i) either $h = g_1$ or $h \triangleright_\ell g_1$ and (ii) ℓ cannot weakly accommodate h alongside \widehat{H}_ℓ^{h} . First, as h has a weakly higher priority than g_1 and $\theta_\ell^{g_1}(G') \neq \infty$ (by inequality (11) and Lemma 1(i)), we can apply Lemma 1(iii) to obtain that $\theta_\ell^{g_1}(G') \leq \theta_\ell^h(G')$. Second, as $G' \subseteq H'$, we have that $\widehat{G}_\ell^h \subseteq \widehat{H}_\ell^h$; therefore, \widehat{H}_ℓ^h is a subset of \widehat{F}_ℓ^h that contains all families in \widehat{G}_ℓ^h and alongside which ℓ cannot weakly accommodate h . By the definition of thresholds (Algorithm 5), it follows that $\theta_\ell^h(G') \leq |\widehat{H}_\ell^h|$. Third, as $H' = G' \cup \{f_i\}$, we have that $|\widehat{H}_\ell^h| \leq |\widehat{G}_\ell^h| + 1$. Fourth, as h has a weakly higher priority than g_1 , we have that $|\widehat{G}_\ell^h| \leq |\widehat{G}_\ell^{g_1}|$. Combining these four observations yields

$$\theta_\ell^{g_1}(G') \leq \theta_\ell^h(G') \leq |\widehat{H}_\ell^h| \leq |\widehat{G}_\ell^h| + 1 \leq |\widehat{G}_\ell^{g_1}| + 1,$$

which contradicts inequality (12).

Sub-sub-case 2.2.2: $\theta_\ell^{g_1}(H') \in \mathbb{Z}_{>0}$. In this case, we can apply Lemma 1(ii) to obtain that $\theta_\ell^{g_1}(H') \geq \theta_\ell^{g_1}(G')$. Moreover, as $H' = G' \cup \{f_i\}$ and $f_i \in \widehat{H}_\ell^{g_1}$ (by the assumption of Sub-case 2.2), $f_i \triangleright_\ell g_1$), we have that $|\widehat{H}_\ell^{g_1}| = |\widehat{G}_\ell^{g_1}| + 1$. Combining these observations with inequality (12) yields

$$\theta_\ell^{g_1}(H') \geq \theta_\ell^{g_1}(G') \geq |G_\ell^{g_1}| + 2 = |H_\ell^{g_1}| + 1,$$

which contradicts inequality (11) and completes the proof of Claim 2.

TKDA is Weakly Envy-free

Towards a contradiction, suppose that μ^{TKDA} is not weakly envy-free. Then, by Definition 10(iii), there exists a family $f' \in F$ matched to a locality $\ell' \in L$ (i.e., $\mu^{TKDA}(f') = \ell'$) such that ℓ' cannot weakly accommodate f' alongside $\widehat{G} = \{g \in F : g \triangleright_{\ell'} f' \text{ and } \ell' \succeq_g \mu^{TKDA}(g)\}$. Let N be the last round of the TKDA algorithm and, for all $i = 1, \dots, N$, denote by $\Pi_{\ell'}^i$ the set of families that propose to ℓ' and by $\theta_{\ell'}^{f'}(\Pi_{\ell'}^i)$ the threshold of family f' for locality ℓ' in Round i .

Consider Round N . By construction, ℓ' does not permanently reject any family and $\mu^{TKDA}(\ell') = \Pi_{\ell'}^N$. If $\theta_{\ell'}^{f'}(\Pi_{\ell'}^N) = \infty$, then locality ℓ' can weakly accommodate f' alongside $\widehat{F}_{\ell'}^{f'}$. This is a contradiction since $\widehat{G} \subseteq \widehat{F}_{\ell'}^{f'}$. Therefore, $\theta_{\ell'}^{f'}(\Pi_{\ell'}^N) \neq \infty$.

Suppose now that, throughout the algorithm, ℓ' does not permanently reject any family that has a higher priority than f' (i.e., ℓ' does not permanently reject any family in $\widehat{F}_{\ell'}^{f'}$). In that case, all families in $\widehat{F}_{\ell'}^{f'}$ are matched to either ℓ' or a more preferred locality so $\widehat{G} = \Pi_{\ell'}^N \cap \widehat{F}_{\ell'}^{f'}$. Since ℓ' does not permanently reject f' , ℓ' can weakly accommodate f' alongside $\widehat{G} = \Pi_{\ell'}^N \cap \widehat{F}_{\ell'}^{f'}$, a contradiction.

We have therefore established that $\theta_{\ell'}^{f'}(\Pi_{\ell'}^N) \neq \infty$ and that locality ℓ' permanently rejects at least one family that has a higher priority than f' . Let g be the highest-priority family that ℓ' permanently rejects (by assumption, $g \triangleright_{\ell'} f'$) and let $i = 1, \dots, N-1$ be the round in which ℓ' permanently rejects g . All families that have a higher priority than g and that propose to ℓ' at some point in the algorithm must also propose to ℓ' in Round N so $(\Pi_{\ell'}^i \cap \widehat{F}_{\ell'}^g) \subseteq (\Pi_{\ell'}^N \cap \widehat{F}_{\ell'}^g)$.

There are two remaining cases to consider: $\theta_{\ell'}^g(\Pi_{\ell'}^i) = 0$ and $\theta_{\ell'}^g(\Pi_{\ell'}^i) \in \mathbb{Z}_{>0}$.

Case 1: $\theta_{\ell'}^g(\Pi_{\ell'}^i) = 0$. By definition, locality ℓ' cannot weakly accommodate g alongside $\Pi_{\ell'}^i \cap \widehat{F}_{\ell'}^g$; therefore ℓ' cannot weakly accommodate g alongside $\Pi_{\ell'}^N \cap \widehat{F}_{\ell'}^g$, meaning that $\theta_{\ell'}^g(\Pi_{\ell'}^N) = 0$. By Lemma 1(iii), $\theta_{\ell'}^{f'}(\Pi_{\ell'}^N) = 0$ so ℓ' permanently rejects f' in Round N , a contradiction.

Case 2: $\theta_{\ell'}^g(\Pi_{\ell'}^i) \in \mathbb{Z}_{>0}$. By assumption, locality ℓ' permanently rejects g in Round i . First, since $(\Pi_{\ell'}^i \cap \widehat{F}_{\ell'}^g) \subseteq (\Pi_{\ell'}^N \cap \widehat{F}_{\ell'}^g)$, Lemma 1(ii) applies and yields

$$\theta_{\ell'}^g(\Pi_{\ell'}^N \cap \widehat{F}_{\ell'}^g) \leq \theta_{\ell'}^g(\Pi_{\ell'}^i \cap \widehat{F}_{\ell'}^g) + |\Pi_{\ell'}^N \cap \widehat{F}_{\ell'}^g| - |\Pi_{\ell'}^i \cap \widehat{F}_{\ell'}^g|. \quad (13)$$

Second, from Algorithm 5, the threshold of a family at a given locality only depends on

families with a higher priority at that locality; therefore

$$\theta_{\ell'}^g(\Pi_{\ell'}^i) = \theta_{\ell'}^g(\Pi_{\ell'}^i \cap \widehat{F}_{\ell'}^g) \quad \text{and} \quad \theta_{\ell'}^g(\Pi_{\ell'}^N) = \theta_{\ell'}^g(\Pi_{\ell'}^N \cap \widehat{F}_{\ell'}^g). \quad (14)$$

Finally, by Lemma 1(iii), we have that

$$\theta_{\ell'}^{f'}(\Pi_{\ell'}^N) \leq \theta_{\ell'}^g(\Pi_{\ell'}^N). \quad (15)$$

Combining (13), (14), and (15), we obtain that

$$\theta_{\ell'}^{f'}(\Pi_{\ell'}^N) \leq \theta_{\ell'}^g(\Pi_{\ell'}^N) \leq \theta_{\ell'}^g(\Pi_{\ell'}^i) + |\Pi_{\ell'}^N \cap \widehat{F}_{\ell'}^g| - |\Pi_{\ell'}^i \cap \widehat{F}_{\ell'}^g|. \quad (16)$$

Finally, the fact that $g \triangleright_{\ell'} f'$ implies that $|\Pi_{\ell'}^N \cap \widehat{F}_{\ell'}^g| \leq |\Pi_{\ell'}^N \cap \widehat{F}_{\ell'}^{f'}|$ and the fact that ℓ' permanently rejects g in Round i implies that $\theta_{\ell'}^g(\Pi_{\ell'}^i) < |\Pi_{\ell'}^i \cap \widehat{F}_{\ell'}^g| + 1$, or, equivalently, that $\theta_{\ell'}^g(\Pi_{\ell'}^i) - |\Pi_{\ell'}^i \cap \widehat{F}_{\ell'}^g| < 1$. Therefore, we have that

$$\theta_{\ell'}^g(\Pi_{\ell'}^i) + |\Pi_{\ell'}^N \cap \widehat{F}_{\ell'}^g| - |\Pi_{\ell'}^i \cap \widehat{F}_{\ell'}^g| < |\Pi_{\ell'}^N \cap \widehat{F}_{\ell'}^{f'}| + 1. \quad (17)$$

Combining (16) and (17) yields $\theta_{\ell'}^{f'}(\Pi_{\ell'}^N) < |\Pi_{\ell'}^N \cap \widehat{F}_{\ell'}^{f'}| + 1$. Therefore, ℓ' permanently rejects f' in Round N , a contradiction. \square

Proof of Lemma 1

Proof of (i) If $\theta_{\ell}^f(G) = \infty$, then ℓ can weakly accommodate f alongside \widehat{F}_{ℓ}^f , which implies $\theta_{\ell}^f(H) = \infty$. The converse is proved analogously.

Proof of (ii) We have assumed that $\theta_{\ell}^f(H) \in \mathbb{Z}_{>0}$. We first show that $\tilde{\theta}_{\ell}^f(H), \tilde{\theta}_{\ell}^f(G), \theta_{\ell}^f(G) \in \mathbb{Z}_{>0}$. If either $\tilde{\theta}_{\ell}^f(H) = \infty$, $\tilde{\theta}_{\ell}^f(G) = \infty$, or $\theta_{\ell}^f(G) = \infty$, ℓ can weakly accommodate f alongside \widehat{F}_{ℓ}^f , which implies that $\theta_{\ell}^f(H) = \infty$, a contradiction. It remains to show that none of $\tilde{\theta}_{\ell}^f(H)$, $\tilde{\theta}_{\ell}^f(G)$, or $\theta_{\ell}^f(G)$ are equal to 0. Since $\tilde{\theta}_{\ell}^f(H) \neq \infty$, $\theta_{\ell}^f(H) \leq \tilde{\theta}_{\ell}^f(H)$ by definition (Algorithm 5); therefore $\tilde{\theta}_{\ell}^f(H) = 0$ implies $\theta_{\ell}^f(H) = 0$, a contradiction. If $\tilde{\theta}_{\ell}^f(G) = 0$, then ℓ cannot weakly accommodate f alongside \widehat{G}_{ℓ}^f . Since $\widehat{G}_{\ell}^f \subseteq \widehat{H}_{\ell}^f$, ℓ cannot weakly accommodate f alongside \widehat{H}_{ℓ}^f so $\tilde{\theta}_{\ell}^f(H) = 0$, a contradiction. If $\theta_{\ell}^f(G) = 0$ and $\tilde{\theta}_{\ell}^f(G) \neq 0$, then there exists $g \in \widehat{F}_{\ell}^f$ such that $\tilde{\theta}_{\ell}^g(G) = 0$. Since $G \subseteq H$, we have that $\tilde{\theta}_{\ell}^g(H) = 0$. Hence, as $g \in \widehat{F}_{\ell}^f$, we have that $\theta_{\ell}^f(H) = 0$, a contradiction.

Having established that $\tilde{\theta}_{\ell}^f(H), \tilde{\theta}_{\ell}^f(G), \theta_{\ell}^f(G) \in \mathbb{Z}_{>0}$, we next show that $\tilde{\theta}_{\ell}^f(G) \leq \tilde{\theta}_{\ell}^f(H)$. As $\tilde{\theta}_{\ell}^f(H) \in \mathbb{Z}_{>0}$, there exists a subset of families \tilde{F} such that $\widehat{H}_{\ell}^f \subseteq \tilde{F} \subseteq \widehat{F}_{\ell}^f$ and $|\tilde{F}| = \tilde{\theta}_{\ell}^f(H)$ alongside which ℓ cannot weakly accommodate f . Since $G \subseteq H$, we have that $\widehat{G}_{\ell}^f \subseteq \widehat{H}_{\ell}^f$ and therefore $\widehat{G}_{\ell}^f \subseteq \widehat{H}_{\ell}^f \subseteq \tilde{F}$. Hence, we obtain that $\tilde{\theta}_{\ell}^f(G) \leq |\tilde{F}| = \tilde{\theta}_{\ell}^f(H)$ as required.

We next show that $\tilde{\theta}_{\ell}^f(H) \leq \tilde{\theta}_{\ell}^f(G) + |\widehat{H}_{\ell}^f| - |\widehat{G}_{\ell}^f|$. As $\tilde{\theta}_{\ell}^f(G) \in \mathbb{Z}_{>0}$, there exists a subset of families \tilde{F} such that $\widehat{G}_{\ell}^f \subseteq \tilde{F} \subseteq \widehat{F}_{\ell}^f$ and $|\tilde{F}| = \tilde{\theta}_{\ell}^f(G)$ alongside which ℓ cannot weakly accommodate f . Therefore, ℓ cannot weakly accommodate f alongside $\tilde{F} \cup \widehat{H}_{\ell}^f$ so $\tilde{\theta}_{\ell}^f(H) \leq |\tilde{F} \cup \widehat{H}_{\ell}^f|$. By construction, $|\tilde{F} \cup \widehat{H}_{\ell}^f| = |\tilde{F}| + |\widehat{H}_{\ell}^f| - |\tilde{F} \cap \widehat{H}_{\ell}^f|$ and $\widehat{G}_{\ell}^f \subseteq (\tilde{F} \cap \widehat{H}_{\ell}^f)$

so $|\widehat{G}_\ell^f| \leq |\widetilde{F} \cap \widehat{H}_\ell^f|$. Combining the preceding results with the fact that $|\widetilde{F}| = \widetilde{\theta}_\ell^f(G)$ yields

$$\widetilde{\theta}_\ell^f(H) \leq |\widetilde{F} \cup \widehat{H}_\ell^f| = |\widetilde{F}| + |\widehat{H}_\ell^f| - |\widetilde{F} \cap \widehat{H}_\ell^f| \leq |\widetilde{F}| + |\widehat{H}_\ell^f| - |\widehat{G}_\ell^f| = \widetilde{\theta}_\ell^f(G) + |\widehat{H}_\ell^f| - |\widehat{G}_\ell^f|.$$

We have now established that $\widetilde{\theta}_\ell^f(G) \leq \widetilde{\theta}_\ell^f(H) \leq \widetilde{\theta}_\ell^f(G) + |\widehat{H}_\ell^f| - |\widehat{G}_\ell^f|$. Since f was chosen arbitrarily, we have

$$\min_{g \in \widehat{F}_\ell^f \cup \{f\}} \widetilde{\theta}_\ell^g(G) \leq \min_{g \in \widehat{F}_\ell^f \cup \{f\}} \widetilde{\theta}_\ell^g(H) \leq \min_{g \in \widehat{F}_\ell^f \cup \{f\}} \left\{ \widetilde{\theta}_\ell^g(G) + |\widehat{H}_\ell^g| - |\widehat{G}_\ell^g| \right\}. \quad (18)$$

Let $h = \arg \min_{g \in \widehat{F}_\ell^f \cup \{f\}} \widetilde{\theta}_\ell^g(G)$. Then, we have that

$$\min_{g \in \widehat{F}_\ell^f \cup \{f\}} \left\{ \widetilde{\theta}_\ell^g(G) + |\widehat{H}_\ell^g| - |\widehat{G}_\ell^g| \right\} \leq \widetilde{\theta}_\ell^h(G) + |\widehat{H}_\ell^h| - |\widehat{G}_\ell^h|. \quad (19)$$

In addition, we have that

$$|\widehat{H}_\ell^h| - |\widehat{G}_\ell^h| = |\widehat{H}_\ell^h \setminus \widehat{G}_\ell^h| \leq |\widehat{H}_\ell^f \setminus \widehat{G}_\ell^f| = |\widehat{H}_\ell^f| - |\widehat{G}_\ell^f|, \quad (20)$$

where the two equalities follow from the fact that $\widehat{G}_\ell^h \subseteq \widehat{H}_\ell^h$ and $\widehat{G}_\ell^f \subseteq \widehat{H}_\ell^f$ while the fact that $(\widehat{H}_\ell^h \setminus \widehat{G}_\ell^h) \subseteq (\widehat{H}_\ell^f \setminus \widehat{G}_\ell^f)$ (as $h \in \widehat{F}_\ell^f$) implies the inequality. Combining (18), (19), and (20), we obtain

$$\min_{g \in \widehat{F}_\ell^f \cup \{f\}} \widetilde{\theta}_\ell^g(G) \leq \min_{g \in \widehat{F}_\ell^f \cup \{f\}} \widetilde{\theta}_\ell^g(H) \leq \widetilde{\theta}_\ell^h(G) + |\widehat{H}_\ell^f| - |\widehat{G}_\ell^f|. \quad (21)$$

Finally, using the fact that $\theta_\ell^f(G), \theta_\ell^f(H) \in \mathbb{Z}_{>0}$ and the definition of θ_ℓ^f from Algorithm 5, we obtain that

$$\theta_\ell^f(G) = \min_{g \in \widehat{F}_\ell^f \cup \{f\}} \widetilde{\theta}_\ell^g(G) = \widetilde{\theta}_\ell^h(G) \quad \text{and} \quad \theta_\ell^f(H) = \min_{g \in \widehat{F}_\ell^f \cup \{f\}} \widetilde{\theta}_\ell^g(H). \quad (22)$$

Combining (21) and (22) yields

$$\theta_\ell^f(G) \leq \theta_\ell^f(H) \leq \theta_\ell^f(G) + |\widehat{H}_\ell^f| - |\widehat{G}_\ell^f|,$$

as required.

Proof of (iii) We have assumed that $g \triangleright_\ell f$ and $\theta_\ell^f(G) \neq \infty$. To prove the first part of the statement, we need to show that $\theta_\ell^g(G) \geq \theta_\ell^f(G)$. If $\theta_\ell^g(G) = \infty$, the result is immediate since $\theta_\ell^f(G)$ is finite. If $\theta_\ell^g(G) \neq \infty$, the definition of thresholds in Algorithm 5 implies that

$$\theta_\ell^f(G) = \min_{h \in \widehat{G}_\ell^f} \widetilde{\theta}_\ell^h \quad \text{and} \quad \theta_\ell^g(G) = \min_{h \in \widehat{G}_\ell^g} \widetilde{\theta}_\ell^h.$$

Then, the fact that $g \triangleright_\ell f$ implies $\widehat{G}_\ell^g \subseteq \widehat{G}_\ell^f$; therefore

$$\theta_\ell^g(G) = \min_{h \in \widehat{G}_\ell^g} \widetilde{\theta}_\ell^h \geq \min_{h \in \widehat{G}_\ell^f} \widetilde{\theta}_\ell^h = \theta_\ell^f(G),$$

as required.

To prove the second part of the statement we need to show that $g \in G \Leftrightarrow g \in C_\ell(G)$ under the additional assumption that $f \in C_\ell(G)$. By definition of a choice function, if $g \in C_\ell(G)$, then $g \in G$. It remains to show that $f \in C_\ell(G)$ and $g \in G$ imply that $g \in C_\ell(G)$.

By definition of a choice function, $g \in C_\ell(G)$ whenever $\theta_\ell^g(G) \geq |\widehat{G}_\ell^g| + 1$ or, equivalently, whenever $\theta_\ell^g(G) - |\widehat{G}_\ell^g| \geq 1$. We have already established in the first part of the statement that $\theta_\ell^g(G) \geq \theta_\ell^f(G)$. Additionally, since $g \triangleright_\ell f$ and $g \in G$, we have that $\widehat{G}_\ell^g \subset \widehat{G}_\ell^f$, and therefore $|\widehat{G}_\ell^g| < |\widehat{G}_\ell^f|$. Taken together, these results imply that

$$\theta_\ell^g(G) - |\widehat{G}_\ell^g| > \theta_\ell^f(G) - |\widehat{G}_\ell^f|. \quad (23)$$

Since $f \in C_\ell(G)$ by assumption, we have that $\theta_\ell^f(G) \geq |\widehat{G}_\ell^f| + 1$ (by definition of a choice function). Combining $\theta_\ell^f(G) - |\widehat{G}_\ell^f| \geq 1$ with (23) yields

$$\theta_\ell^g(G) - |\widehat{G}_\ell^g| > \theta_\ell^f(G) - |\widehat{G}_\ell^f| \geq 1.$$

We can conclude that $\theta_\ell^g(G) > |\widehat{G}_\ell^g| + 1$. By definition of a choice function, we have that $g \in C_\ell(G)$, which completes the proof. \square

Proof of Proposition 8

We denote by d the unique dimension (i.e., $D = \{d\}$), and consider each mechanism separately.

Proof that the KDA mechanism is envy-free Let N be the number of rounds after which the KDA algorithm ends and let μ be the outcome of the KDA algorithm. We use a piece of notation that was introduced in the proof of Theorem 5. For any family $f \in F$, any locality $\ell \in L$, and any Round $i = 1, \dots, N$, $\widehat{R}_f^i(\ell)$ denotes the set of families that have a higher priority than f at ℓ and propose in Round i either to ℓ or to a less-preferred locality.

Towards a contradiction, suppose that μ is not envy-free. Then, by definition, there exist two families f and f' such that f envies f' . Letting $\ell' = \mu(f')$, we have that $\ell' \succ_f \mu(f)$ and $f \triangleright_{\ell'} f'$. By construction, there exists a Round i in which ℓ' rejects f ; hence ℓ' cannot weakly accommodate f alongside $\widehat{R}_f^i(\ell')$. Again by construction, we have that $\widehat{R}_f^i(\ell') \subseteq \widehat{R}_f^N(\ell')$ so we conclude that ℓ' cannot weakly accommodate f alongside $\widehat{R}_f^N(\ell')$.

As d is the only dimension, it must be that $\nu_d^{f'} > 0$, and therefore the fact that ℓ' cannot weakly accommodate f alongside $\widehat{R}_f^N(\ell')$ implies by definition that $\nu_d^f + \sum_{g \in \widehat{R}_f^N(\ell')} \nu_d^g > \kappa_d^{\ell'}$. As $f \triangleright f'$, it follows that $\nu_d^{f'} + \sum_{g \in \widehat{R}_{f'}^N(\ell')} \nu_d^g > \kappa_d^{\ell'}$. Therefore, ℓ' cannot weakly accommodate f' alongside $\widehat{R}_{f'}^N(\ell')$ and we conclude that, in Round N , either f' does not propose to ℓ'

or ℓ' permanently rejects f' ; either conclusion contradicts the assumption that $\ell' = \mu(f')$, completing the proof.

Proof that the TKDA mechanism is envy-free Let N be the number of rounds after which the TKDA algorithm ends and let μ be the outcome of the TKDA algorithm. We use two pieces of notation that were introduced in the proof of Theorem 6. First, for any Round $i = 1, \dots, N$ and any locality $\ell \in L$, Π_ℓ^i denotes set of families that propose to locality ℓ in Round i . Second, for any family $f \in F$, any locality $\ell \in L$, and any set of families $G \subseteq F$, $\theta_\ell^f(G)$ denotes f 's threshold for ℓ (calculated by Algorithm 5) when the families in G propose to ℓ . Moreover, we will invoke Lemma 1, which was also introduced in the proof of Theorem 6.

Towards a contradiction, suppose that μ is not envy-free. Then, by definition, there exist two families f and f' such that f envies f' . Without loss of generality, let f be the highest-priority family among those that envy f' . Formally, letting $\ell' = \mu(f')$, we have that $\ell' \succ_f \mu(f)$, that $f \triangleright_{\ell'} f'$, and that for any $\tilde{f} \neq f$ such that $\ell' \succ_{\tilde{f}} \mu(\tilde{f})$, $f \triangleright_{\ell'} \tilde{f}$. As $\ell' \succ_f \mu(f)$, by construction there exists a Round i in which ℓ' permanently rejects f so we have that

$$\theta_{\ell'}^f(\Pi_{\ell'}^i) < |\Pi_{\ell'}^i \cap \widehat{F}_{\ell'}^f| + 1. \quad (24)$$

The next step consists of showing that the following inequality holds:

$$\theta_{\ell'}^f(\Pi_{\ell'}^N) < |\Pi_{\ell'}^N \cap \widehat{F}_{\ell'}^f| + 1. \quad (25)$$

If $\theta_{\ell'}^f(\Pi_{\ell'}^N) = \infty$, then by Lemma 1(i), $\theta_{\ell'}^f(\Pi_{\ell'}^i) = \infty$, which contradicts (24) and, if $\theta_{\ell'}^f(\Pi_{\ell'}^N) = 0$, (25) holds trivially. Therefore, it remains to show that (25) holds in the case where $\theta_{\ell'}^f(\Pi_{\ell'}^N) \in \mathbb{Z}_{>0}$. As f has the highest-priority among the families that envy f' , by construction f is also the highest-priority family that ℓ' permanently rejects throughout the TKDA algorithm; hence we have that $\Pi_{\ell'}^i \cap \widehat{F}_{\ell'}^f \subseteq \Pi_{\ell'}^N \cap \widehat{F}_{\ell'}^f$. By Lemma 1(ii), it follows that

$$\theta_{\ell'}^f(\Pi_{\ell'}^N \cap \widehat{F}_{\ell'}^f) \leq \theta_{\ell'}^f(\Pi_{\ell'}^i \cap \widehat{F}_{\ell'}^f) + |\Pi_{\ell'}^N \cap \widehat{F}_{\ell'}^f| - |\Pi_{\ell'}^i \cap \widehat{F}_{\ell'}^f|.$$

By construction (see Algorithm 5), f 's threshold at ℓ' only depends on higher-priority families so $\theta_{\ell'}^f(\Pi_{\ell'}^N \cap \widehat{F}_{\ell'}^f) = \theta_{\ell'}^f(\Pi_{\ell'}^N)$ and $\theta_{\ell'}^f(\Pi_{\ell'}^i \cap \widehat{F}_{\ell'}^f) = \theta_{\ell'}^f(\Pi_{\ell'}^i)$; hence we have that

$$\theta_{\ell'}^f(\Pi_{\ell'}^N) \leq \theta_{\ell'}^f(\Pi_{\ell'}^i) + |\Pi_{\ell'}^N \cap \widehat{F}_{\ell'}^f| - |\Pi_{\ell'}^i \cap \widehat{F}_{\ell'}^f|,$$

which, combined with (24), implies (25).

Having established that inequality (25) holds, we now use it to reach the contradiction. We consider separately two cases: $\theta_{\ell'}^{f'}(\Pi_{\ell'}^N) \neq \infty$ and $\theta_{\ell'}^{f'}(\Pi_{\ell'}^N) = \infty$.

Case 1: $\theta_{\ell'}^{f'}(\Pi_{\ell'}^N) \neq \infty$. As $f \triangleright_{\ell'} f'$, we have that $\Pi_{\ell'}^N \cap \widehat{F}_{\ell'}^f \subseteq \Pi_{\ell'}^N \cap \widehat{F}_{\ell'}^{f'}$ and, by Lemma 1(iii), $\theta_{\ell'}^{f'}(\Pi_{\ell'}^N) \leq \theta_{\ell'}^f(\Pi_{\ell'}^N)$. Then, (25) implies that $\theta_{\ell'}^{f'}(\Pi_{\ell'}^N) < |\Pi_{\ell'}^N \cap \widehat{F}_{\ell'}^{f'}| + 1$. It follows that, in Round N , either f' does not propose to ℓ' or ℓ' permanently rejects f' , both of which contradict the assumption that $\ell' = \mu(f')$.

Case 2: $\theta_{\ell'}^{f'}(\Pi_{\ell'}^N) = \infty$. By definition (see Algorithm 5), the case assumption implies that ℓ' can weakly accommodate f' alongside $\widehat{F}_{\ell'}^{f'}$. As d is the only dimension, we have that

$\nu_d^{f'} > 0$ so, by definition: $\nu_d^{f'} + \sum_{g \in \widehat{F}_{\ell'}^{f'}} \nu_d^g \leq \kappa_d^{\ell'}$. It follows that $\nu_d^f + \sum_{g \in \widehat{F}_{\ell'}^f} \nu_d^g < \kappa_d^{\ell'}$; hence ℓ' can (weakly) accommodate f alongside $\widehat{F}_{\ell'}^f$. By definition (see Algorithm 5), we conclude that $\theta_{\ell'}^f(\Pi_{\ell'}^N) = \infty$, which contradicts (25). \square

Online Appendix

A Efficiency of the TKDA algorithm

A.1 Lower bound on the efficiency of KDA and TKDA mechanisms

We now show the lower bound on the efficiency of the KDA and TKDA mechanisms.

Proposition 9. *The KDA and TKDA mechanisms match at least one family to a locality in $L \setminus \{\emptyset\}$. If all families can be accommodated on their own at all localities, then the KDA and TKDA mechanisms match at least $\min\{|F|, |L| - 1\}$ families to localities in $L \setminus \{\emptyset\}$.*

The intuition for Proposition 9 is as follows. First, if a family f that can be accommodated at some locality ℓ is matched to the null locality, then f has been permanently rejected by ℓ . Therefore, there exists a family f' with a higher priority than f at ℓ that (i) can be accommodated on its own at ℓ and (ii) proposes to ℓ in some round of the KDA or TKDA algorithm. In turn, family f' would only be matched to the null locality if there is a yet another higher priority family f'' and, by induction, it is not possible that all families be matched to the null locality. Second, consider the case where all families can be accommodated on their own at all localities. If fewer than $|L| - 1$ families are matched to non-null localities, then either there are fewer than $|L| - 1$ families in the market or at least one family f is matched to the null locality while a non-null locality ℓ is not matched to any family. This yields a contradiction since f proposes to and is permanently rejected by ℓ in some round of the KDA or TKDA algorithms.

A.2 TKDA with Clinching algorithm

In this section, we present a modification of the TKDA algorithm that improves its efficiency without affecting its properties. The TKDA with Clinching (TKDAC) algorithm (Algorithm 7) starts with a Clinching Round that creates a new priority profile. The TKDA algorithm is then run using the new priority profile.

Proposition 10. *The TKDAC mechanism is strategy-proof and weakly envy-free. Moreover, $\mu^{TKDAC} \succeq \mu^{TKDA}$.*

The idea of the TKDAC algorithm is to identify family-locality pairs that will necessarily be matched together by the TKDA algorithm. If locality ℓ is family f 's first preference and ℓ can weakly accommodate f alongside all higher-priority families, then $\mu^{TKDA}(f) = \ell$: f proposes to ℓ and ℓ tentatively accepts f 's proposal throughout the algorithm because f 's threshold at ℓ is ∞ . In the Clinching Round of the TKDAC algorithm, family f clinches locality ℓ . Then we construct a new priority profile in which f moves to the bottom of the priority list of every locality ℓ' such that $\ell \succ_f \ell'$. The change in the priority profile does not affect f 's match in the TKDA algorithm, since f will be matched to ℓ , no matter what f 's priority for less-preferred localities. However, the change in the priority profile may positively affect families that propose to f 's less-preferred localities since their thresholds at these localities are no longer affected by f .

	Round 1			Round 2			Round 3			Round 4		
$f_1 \rightarrow$	ℓ_1	[1]	✓	ℓ_1	[0]	✗	ℓ_2	[∞]	✓	ℓ_2	[∞]	✓
$f_2 \rightarrow$	ℓ_2	[1]	✓	ℓ_2	[1]	✓	ℓ_2	[0]	✗	ℓ_1	[∞]	✓
$f_3 \rightarrow$	ℓ_2	[1]	✗	ℓ_1	[∞]	✓	ℓ_1	[∞]	✓	ℓ_1	[∞]	✓

Table 9: TKDAC* algorithm applied to Example 4.

To illustrate how clinching can improve families' welfare, consider a locality ℓ_1 and three families f_1 , f_2 , and f_3 such that $\triangleright_{\ell_1} : f_1, f_2, f_3, \dots$. There is only one dimension d_1 and the sizes and capacities are

$$d_1 \begin{pmatrix} f_1 & f_2 & f_3 & \ell_1 \\ 2 & 1 & 1 & 2 \end{pmatrix}.$$

Consider what happens if f_2 and f_3 both propose to ℓ_1 in some round of the TKDA algorithm. The thresholds are $\theta_{\ell_1}^{f_1} = \infty$ (since f_1 has the highest priority for ℓ_1 and ℓ_1 can accommodate f_1 on its own) and $\theta_{\ell_1}^{f_2} = \theta_{\ell_1}^{f_3} = 1$ (since ℓ_1 can (weakly) accommodate each of f_2 or f_3 on its own but not alongside f_1). Therefore, ℓ_1 tentatively accepts f_2 's proposal (since f_2 's priority rank among proposing families is $1 = \theta_{\ell_1}^{f_2}$) and permanently rejects f_3 's proposal (since f_3 's priority rank among proposing families is $2 > 1 = \theta_{\ell_1}^{f_3}$). Observe that ℓ_1 permanently rejects f_3 even though ℓ_1 can (weakly) accommodate both proposing families.

Suppose, however, that f_1 's first preference is another locality ℓ_2 and that ℓ_2 can (weakly) accommodate f_1 alongside all higher-priority families. Then, the clinching round identifies this pair and f_1 clinches ℓ_2 . The Clinching Round produces a priority profile $\tilde{\triangleright}$ such that $f_2 \tilde{\triangleright}_{\ell_1} f_3 \tilde{\triangleright}_{\ell_1} f_1$. If the TKDA algorithm is run with the new priority profile $\tilde{\triangleright}$ and f_2 and f_3 both propose to ℓ_1 , we have that $\theta_{\ell_1}^{f_2} = \theta_{\ell_1}^{f_3} = \infty$ since f_2 and f_3 now have the two highest priorities for ℓ_1 and therefore ℓ_1 can accommodate both f_2 and f_3 . As a consequence, following the Clinching Round, ℓ_1 no longer permanently rejects f_3 resulting into a Pareto improvement over the original TKDA mechanism.

The TKDAC algorithm also allows families to clinch localities that are not their first preferences. In Step 1 of the Clinching Round, locality ℓ rejects family f if ℓ cannot accommodate f on its own (as no clinches have occurred before the start of Step 1). If ℓ rejects f and f receives a proposal from its second-preference locality ℓ' , i.e., if ℓ' can weakly accommodate f alongside all higher-priority families, then it can be established that the TKDA algorithm would match f to ℓ' . Therefore, f clinches ℓ' , i.e., f goes to the bottom of the priority list of all of f 's less-preferred localities. In Step 2, these localities may propose to new families as a result. In addition, ℓ' now rejects every family with a lower priority than f at ℓ' that ℓ' cannot weakly accommodate alongside f . The Clinching Round continues until there is a step in which no family clinches any locality. Online Appendix B.4 provides an example of the TKDAC algorithm.

A.3 Why clinching any proposing locality affects incentives for truth-telling

As clinching improves efficiency, one might consider allowing families to clinch any locality that proposes to them (whether or not they have been rejected by all of their more preferred localities). Suppose we allowed family f to clinch a locality ℓ as long as ℓ can weakly

Algorithm 7: TKDA WITH CLINCHES (TKDAC)

Clinching Round

Step 0

No locality rejects or proposes to any family and no family clinches any locality.

Set $\tilde{\succ}^1 = \triangleright$ and continue to Step 1.

Step $j \geq 1$

(a) Locality ℓ rejects family f if ℓ cannot weakly accommodate f alongside the families that (i) clinched ℓ in Step $j - 1$ and (ii) are higher than f on $\tilde{\succ}_\ell^j$.

(b) Locality ℓ proposes to family f if ℓ can weakly accommodate f alongside all families that are higher than f on $\tilde{\succ}_\ell^j$.

(c) Family f clinches locality ℓ if (i) ℓ proposed to f in part (b) and (ii) ℓ is f 's most preferred locality that did not reject f in part (a).

(d) If at least one clinch occurred in part (c) that did not occur in Step $j - 1$, continue to part (e). Otherwise, set $\tilde{\succ} = \tilde{\succ}^j$ and continue to the TKDA algorithm.

(e) Construct $\tilde{\succ}^{j+1}$ as follows, then continue to Step $j + 1$. For every $\ell \in L$ and every $f, f' \in F$ with $f \triangleright_\ell f'$,

- $f' \tilde{\succ}_\ell^{j+1} f$ if, in part (c), (i) f clinched a locality that f strictly prefers to ℓ and (ii) f' did not clinch a locality that f' strictly prefers to ℓ ;
- $f \tilde{\succ}_\ell^{j+1} f'$ otherwise.

TKDA

Run the TKDA algorithm with the priority profile $\tilde{\succ}$.

accommodate f alongside all higher-priority families, even if there are other localities that f prefers and that have not rejected f .

Formally, consider the following modification to the Clinching Round in Algorithm 7: remove condition (ii) of part (c) in every Step j . We call TKDAC* algorithm the TKDAC algorithm with the modified Clinching Round. Therefore, the TKDAC* algorithms differs from the TKDAC algorithm by the fact that, in the Clinching Round, a family f clinches *any* locality that proposes to f .

Unfortunately, the TKDAC* algorithm is not strategy-proof, as the following example shows.

Example 4. There are three families, three localities, and one dimension. The preferences, priorities, family sizes, and locality capacities are

$$\begin{array}{lll}
\succ_{f_1}: \ell_1, \ell_2, \ell_3 & \succ_{f_2}: \ell_2, \ell_1, \ell_3 & \succ_{f_3}: \ell_1, \ell_2, \ell_3 \\
\triangleright_{\ell_1}: f_2, f_3, f_1 & \triangleright_{\ell_2}: f_1, f_2, f_3 & \triangleright_{\ell_3}: f_1, f_2, f_3 \\
& \begin{array}{ccc} f_1 & f_2 & f_3 \end{array} & & \begin{array}{ccc} \ell_1 & \ell_2 & \ell_3 \end{array} \\
\nu = d_1 \begin{pmatrix} 2 & 1 & 1 \end{pmatrix} & & \kappa = d_1 \begin{pmatrix} 2 & 2 & 2 \end{pmatrix}.
\end{array}$$

Suppose that all families report truthfully. In the modified Clinching Round, every family receives a proposal from its second-preference locality. Since ℓ_3 is every families' third- and last-preference locality, the modified Clinching Round does not affect the localities' priorities. As a result, the priority profile $\tilde{\triangleright}$ that is used in the TKDA algorithm is the same as the original priority profile \triangleright .⁴²

The TKDA algorithm—summarized in Table 4—yields the following matching:

$$\begin{pmatrix} f_1 & f_2 & f_3 \\ \ell_2 & \ell_1 & \ell_1 \end{pmatrix}.$$

Suppose now that family f_1 reports $\succ'_{f_1}: \ell_1, \ell_3, \ell_2$ while the other two families report their true preferences. In the modified Clinching Round, f_1 clinches ℓ_3 and drops to the bottom of ℓ_2 's priority list; therefore, we have a new priority list $\tilde{\triangleright}_{\ell_2}: f_2, f_3, f_1$.⁴³ Families f_2 and f_3 clinch locality ℓ_1 , but this proposal is inconsequential since f_3 already has the highest priority for ℓ_3 . Therefore, the modified Clinching Round does not modify the priorities of ℓ_1 and ℓ_3 , i.e., $\tilde{\triangleright}_{\ell_1} = \triangleright_{\ell_1}$ and $\tilde{\triangleright}_{\ell_3} = \triangleright_{\ell_3}$.

The TKDA algorithm is then run with the priority profile $\tilde{\triangleright}$. In the first round, f_1 proposes to ℓ_1 and is tentatively accepted as $\theta_{\ell_1}^{f_1} = 1$ and no other family proposes to ℓ_1 . Families f_2 and f_3 both propose to ℓ_2 . Since f_2 and f_3 both now have a higher priority than f_1 and ℓ_2 can accommodate them together and ℓ_2 tentatively accepts both families (as $\theta_{\ell_2}^{f_2} = \theta_{\ell_2}^{f_3} = \infty$). The TKDA algorithm ends and yields the matching

$$\begin{pmatrix} f_1 & f_2 & f_3 \\ \ell_1 & \ell_2 & \ell_2 \end{pmatrix}.$$

Clearly, f_1 's manipulation has been successful since f_1 is now matched to its first-preference locality ℓ_1 instead of its second-preferences locality ℓ_2 . The reason why f can successfully manipulate is that by clinching ℓ_3 , f_1 allows f_3 to be tentatively accepted by ℓ_2 ; as a result, f_3 does not compete with f_1 for ℓ_1 .

Note that the TKDAC algorithm precludes f_1 's manipulation opportunity in Example 4 because the TKDAC algorithm only allows f_1 to clinch a ℓ_3 when it has been established that the TKDA algorithm will not match f_1 to any locality that f_1 prefers to ℓ_3 (i.e., ℓ_1).

⁴²Family f_1 also receives a proposal from ℓ_3 , but this proposal is inconsequential since ℓ_3 is f_1 's least-preferred locality.

⁴³Family f_1 also receives a proposal from ℓ_2 , but this proposal is inconsequential since ℓ_2 is f_1 's reported least-preferred locality.

A.4 Proofs

Proof of Proposition 9

We first show that the TKDA mechanism matches at least one family to a locality. Recall that we assume throughout that for every family $f \in F$, there exists a locality $\ell \in L \setminus \{\emptyset\}$ such that ℓ can accommodate f on its own (see Section 3). Without loss of generality, let f be the highest-priority family among those that ℓ can accommodate on their own. Recall that we assume throughout that localities prioritize families they can accommodate on their own over families that they cannot accommodate (see Section 3) so f has the highest priority at ℓ among *all* families. Therefore, $\theta_\ell^f = \infty$ throughout the TKDA algorithm, which means that ℓ does not permanently reject f in any round of the TKDA algorithm. If f proposes to ℓ , $\mu^{TKDA}(f) = \ell \succ_f \emptyset$; otherwise, f does not propose to ℓ , so $\mu^{TKDA}(f) \succ_f \ell \succ_f \emptyset$. In both cases, f is matched to a locality that is not the null.

We next show that if all families can be accommodated on their own at all localities then the TKDA mechanism matches at least $\min\{|F|, |L| - 1\}$ families to a locality other than the null. Towards a contradiction, suppose that $|F \setminus \mu^{TKDA}(\emptyset)| < \min\{|F|, |L| - 1\}$. Then, $\mu^{TKDA}(\emptyset) \neq \emptyset$ and there exists $\ell \in L \setminus \{\emptyset\}$ such that $\mu^{TKDA}(\ell) = \emptyset$. Since we assume that \emptyset is every family's last preference and $\mu^{TKDA}(\ell) = \emptyset$, ℓ receives at least one proposal. Let f be the highest-priority family among those that propose to ℓ at least once in the TKDA algorithm. Since the hypothesis states that ℓ can accommodate f on its own, $\theta_\ell^f \geq 1$ throughout the TKDA algorithm. In every round in which f proposes to ℓ , f has the highest priority among proposing families so ℓ never permanently rejects f . This means that $f \in \mu^{TKDA}(\ell)$, which contradicts $\mu^{TKDA}(\ell) = \emptyset$.

Since the KDA mechanism is family optimal, we have that $\mu^{KDA} \succeq \mu^{TKDA}$. Moreover, \emptyset is every family's last preference. Hence, the result also holds for the KDA mechanism. \square

Proof of Proposition 10

First notice that the TKDAC algorithm simply runs the TKDA algorithm with a different priority profile. Since the matching produced by the TKDA algorithm does not violate any priority constraint (see the proof of Theorem 6), the same is true of the TKDAC algorithm, i.e., for every $\ell \in L$, ℓ can accommodate $\mu^{TKDAC}(\ell)$.

We next introduce some notation, which we use throughout the proof. Consider the Clinching Round and let N be its total number of steps. We use throughout the convention that $\tilde{\triangleright}^0 = \triangleright$. For every Step $j = 0, 1, \dots, N$, every $f \in F$, and every $\ell \in L$, let

- $\widehat{F}_\ell^f(\tilde{\triangleright}^j)$ be the set of families that are higher than f on $\tilde{\triangleright}_\ell^j$;
- Δ_f^j be the set of localities that reject f in Step j ;
- Γ_f^j be the set of localities that propose to f in Step j ; and
- Θ_ℓ^j be the set of families that clinch ℓ in Step j .

Note that $\widehat{F}_\ell^f(\tilde{\triangleright}^0) = \widehat{F}_\ell^f(\tilde{\triangleright}^1) = \widehat{F}_\ell^f(\triangleright) = \widehat{F}_\ell^f$ and $\widehat{F}_\ell^f(\tilde{\triangleright}^N) = \widehat{F}_\ell^f(\tilde{\triangleright})$. Since no rejections, proposals, or clinches occur in Step 0, we also have $\Delta_f^0 = \Gamma_f^0 = \Theta_\ell^0 = \emptyset$.

The following two lemmata are key to our analysis. Their proofs can be found directly after the current proof.

Lemma 2. *For every Step $j = 1, \dots, N$, every family $f \in F$, and every locality $\ell \in L$:*

- (i) *If $f \notin \Theta_{\ell'}^{j-1}$ for all $\ell' \in L$ such that $\ell' \succ_f \ell$, then $\widehat{F}_{\ell}^f(\bar{\succ}^j) \subseteq \widehat{F}_{\ell}^f(\bar{\succ}^{j-1})$;*
- (ii) $\Delta_f^{j-1} \subseteq \Delta_f^j$;
- (iii) *If $\ell \in \Gamma_f^{j-1} \setminus \Gamma_f^j$, then there exists $\ell' \in L$ such that $\ell' \succ_f \ell$ and $f \in \Theta_{\ell'}^{j-1}$; and*
- (iv) $\Theta_{\ell}^{j-1} \subseteq \Theta_{\ell}^j$.

Part (i) of Lemma 2 says that if family f has not clinched a locality that f prefer to ℓ , then the set of families who are higher than f at ℓ shrinks throughout the Clinching Round. Part (ii) says that the set of localities that are rejecting f grows throughout the Clinching Round. Part (iii) says that if ℓ stops proposing to f , then f has clinched a more preferred locality. Part (iv) says that the set of families that have clinched ℓ grows throughout the Clinching Round.

Lemma 3. *For every Step $j = 1, \dots, N$, every family $f \in F$, and every locality $\ell \in L$:*

- (i) *If $\ell \in \Delta_f^j$, then $\mu^{TKDAC}(f) \neq \ell$;*
- (ii) *If $\ell \in \Gamma_f^j$, then $\mu^{TKDAC}(f) \succeq_f \ell$; and*
- (iii) *If $f \in \Theta_{\ell}^j$, then $\mu^{TKDAC}(f) = \ell$.*

Part (i) of Lemma 3 says that if a locality ℓ rejects a family f in the Clinching Round, then the TKDAC algorithm will not match f to ℓ . Part (ii) says that if ℓ proposes to f in the Clinching Round, then f will be matched to ℓ or a more preferred locality under the TKDAC algorithm. Part (iii) says that if f clinches ℓ in the Clinching Round, then the TKDAC algorithm will match f to ℓ .

We now use Lemmata 2 and 3 to show that the TKDAC mechanism is strategy-proof and weakly envy-free and that $\mu^{TKDAC} \succeq \mu^{TKDA}$.

TKDAC is Strategy-proof

We consider a family f and a report \succ'_f . We need to show that

$$\varphi^{TKDAC}(\succ)(f) \succeq \varphi^{TKDAC}(\succ'_f, \succ_{-f})(f).$$

We use our usual notation— N , $\bar{\succ}^j$, Δ_f^j , Γ_f^j , and Θ_{ℓ}^j —for the TKDAC algorithm run with the preference profile \succ , i.e., when f reports truthfully. We denote by \bar{N} , $\bar{\succ}^j$, $\bar{\Delta}_f^j$, $\bar{\Gamma}_f^j$, and $\bar{\Theta}_{\ell}^j$ the counterparts in the TKDAC algorithm run with the preference profile (\succ'_f, \succ_{-f}) , i.e., when f misreports its preferences.

Consider the Clinching Round when the preference profile is \succ . If f clinches a locality, let Step m be the first step in which f clinches the locality; formally, $f \notin \cup_{\ell \in L} \{\Theta_{\ell}^j\}$ for

all $j = 1, \dots, m-1$ and $f \in \cup_{\ell \in L} \{\Theta_\ell^j\}$ for all $j = m, \dots, N$. (By Lemma 2(iv), once a family clinches a locality, it continues to clinch the same locality in all remaining steps so m is well defined.) If f does not clinch any locality, let $m = \infty$; formally, $m = \infty$ whenever $f \notin \cup_{\ell \in L} \{\Theta_\ell^N\}$. We define \bar{m} analogously for the preference profile (\succ'_f, \succ_{-f}) . The following lemma guarantees that the Clinching Round is unaffected by f 's report until f clinches a locality.

Lemma 4. *Let $q = \min\{m, \bar{m}, N, \bar{N}\}$. For every $j = 1, \dots, q$, every $g \in F$, and every $\ell \in L$,*

$$\Theta_\ell^j \setminus \{f\} = \bar{\Theta}_\ell^j \setminus \{f\}, \quad \Delta_g^j = \bar{\Delta}_g^j, \quad \Gamma_g^j = \bar{\Gamma}_g^j, \quad \text{and} \quad \tilde{\succ}^j = \bar{\succ}^j.$$

Moreover, if $j < \min\{m, \bar{m}\}$, then $\Theta_\ell^j = \bar{\Theta}_\ell^j$ for all $\ell \in L$.

(The proof of Lemma 4 follows after the current proof.) There are four cases to consider.

Case 1: $m = \infty$ and $\bar{m} = \infty$. In this case, f does not clinch any locality in the Clinching Round, irrespective of whether f reports \succ_f or \succ'_f . Then, $q = \min\{N, \bar{N}\} < \min\{m, \bar{m}\}$. If $N \leq \bar{N}$, then $q = N$ and $\Theta_\ell^{N-1} = \Theta_\ell^N$ for all $\ell \in L$ since, by construction, the Clinching Round ends when the same clinches occur in two consecutive steps. By Lemma 4, $\Theta_\ell^{N-1} = \bar{\Theta}_\ell^{N-1}$ and $\Theta_\ell^N = \bar{\Theta}_\ell^N$ for all $\ell \in L$; therefore, $\bar{\Theta}_\ell^{N-1} = \bar{\Theta}_\ell^N$ for all $\ell \in L$, which means that $N = \bar{N}$. Then, since $\tilde{\succ}^N = \bar{\succ}^N$ by Lemma 4, we conclude that $\tilde{\succ} = \bar{\succ}$. We have established that, whether f reports \succ_f or \succ'_f , the Clinching Round ends in the same step and produces the same adjusted priority profile. Then, $\varphi^{TKDAC}(\succ)(f)$ is the matching produced by the TKDA algorithm when the preference and priority profiles are \succ and $\tilde{\succ}$ respectively while $\varphi^{TKDAC}(\succ'_f, \succ_{-f})(f)$ is the matching produced by the TKDA algorithm when the preference and priority profiles are (\succ'_f, \succ_{-f}) and $\bar{\succ} = \tilde{\succ}$ respectively. Since the TKDA algorithm is strategy-proof (Theorem 6), we conclude that $\varphi^{TKDAC}(\succ)(f) \succeq_f \varphi^{TKDAC}(\succ'_f, \succ_{-f})(f)$, as required. Analogous reasoning yields the same result for the case where $N \geq \bar{N}$.

Case 2: $m \leq \min\{\bar{m}, N\}$. In this case, if f reports truthfully, then f clinches a locality in Step m of the Clinching Round and, if f reports \succ'_f , then f either clinches a locality in Step $\bar{m} \geq m$ or does not clinch any locality. Since $m \leq \min\{\bar{m}, N\}$, $q = \min\{m, \bar{N}\}$. Towards a contradiction, suppose that $m > \bar{N}$. Then, $q = \bar{N} < \min\{m, \bar{m}\}$. As the Clinching Round ends whenever the same clinches occur in two consecutive rounds, we have that $\Theta_\ell^{\bar{N}-1} = \Theta_\ell^{\bar{N}}$ for all $\ell \in L$. Moreover, Lemma 4 implies that $\Theta_\ell^{\bar{N}-1} = \bar{\Theta}_\ell^{\bar{N}-1}$ and $\Theta_\ell^{\bar{N}} = \bar{\Theta}_\ell^{\bar{N}}$ for all $\ell \in L$. It follows that $\bar{\Theta}_\ell^{\bar{N}-1} = \bar{\Theta}_\ell^{\bar{N}}$, so $N = \bar{N}$, which contradicts our assumption that $\bar{N} < m \leq N$.

We have established that $m \leq \bar{N}$; hence $q = m$. When f reports truthfully, f clinches a locality denoted by ℓ in Step m of the Clinching Round, i.e., $f \in \Theta_\ell^m$. By Lemma 3(iii), we have that $\varphi^{TKDAC}(\succ)(f) = \ell$. By construction (Step m (c) of the Clinching Round), $f \in \Theta_\ell^m$ implies that, for all $\ell' \in L$ with $\ell' \succ_f \ell$, $\ell' \in \Delta_f^m$. By Lemma 4, $\Delta_f^m = \bar{\Delta}_f^m$; therefore $\ell' \in \bar{\Delta}_f^m$ for all $\ell' \in L$ such that $\ell' \succ_f \ell$. By Lemma 3(i), it follows that $\varphi^{TKDAC}(\succ'_f, \succ_{-f})(f) \neq \ell'$ for all $\ell' \in L$ such that $\ell' \succ_f \ell$; therefore $\ell \succeq_f \varphi^{TKDAC}(\succ'_f, \succ_{-f})(f)$, as required.

Case 3: $\bar{m} \leq \min\{m, \bar{N}\}$. In this case, if f reports \succ'_f , then f clinches a locality in Step \bar{m} of the Clinching Round and, if f reports truthfully, then f either clinches a locality in Step $m \geq \bar{m}$ or does not clinch any locality. Since $\bar{m} \leq \min\{m, \bar{N}\}$, $q = \min\{\bar{m}, N\}$. Using

analogous reasoning to Case 2, we establish that $\bar{m} \leq N$; hence $q = \bar{m}$. When f reports \succ'_f , f clinches a locality denoted by ℓ in Step \bar{m} of the Clinching Round, i.e., $f \in \bar{\Theta}_\ell^{\bar{m}}$. By Lemma 3(iii), we have that $\varphi^{TKDAC}(\succ'_f, \succ_{-f}) = \ell$. By construction (Step $\bar{m}(c)$ of the Clinching Round), $f \in \bar{\Theta}_\ell^{\bar{m}}$ implies that $\ell \in \bar{\Gamma}_f^{\bar{m}}$. By Lemma 4, $\Gamma_f^{\bar{m}} = \bar{\Gamma}_f^{\bar{m}}$ so $\ell \in \Gamma_f^{\bar{m}}$; therefore, by Lemma 3(ii), $\varphi^{TKDAC}(\succ)(f) \succeq_f \ell$, as required.

TKDAC is Weakly Envy-free

For ease of notation, let $\ell = \mu^{TKDAC}(f)$.

Claim 3. *For any family $f \in F$,*

$$\{g \in F : g \triangleright_\ell f \text{ and } \ell \succeq_g \mu^{TKDAC}(g)\} \subseteq \{g \in F : g \tilde{\triangleright}_\ell f \text{ and } \ell \succeq_g \mu^{TKDAC}(g)\}.$$

Consider a family $f \in F$ and suppose, towards a contradiction, that there exists another family $h \in F$ such that

$$h \in \{g \in F : g \triangleright_\ell f \text{ and } \ell \succeq_g \mu^{TKDAC}(g)\} \setminus \{g \in F : g \tilde{\triangleright}_\ell f \text{ and } \ell \succeq_g \mu^{TKDAC}(g)\}.$$

By assumption, $h \triangleright_\ell f$ and $f \tilde{\triangleright}_\ell h$; equivalently, $f \in \widehat{F}_\ell^h(\tilde{\triangleright}) \setminus \widehat{F}_\ell^h$. Therefore, there exists a Step $j = 1, \dots, N$ of the Clinching Round such that $f \in \widehat{F}_\ell^h(\tilde{\triangleright}^j) \setminus \widehat{F}_\ell^h(\tilde{\triangleright}^{j-1})$. By the contrapositive of Lemma 2(i), it follows that $h \in \Theta_{\ell'}^{j-1}$ for some $\ell' \in L$ such that $\ell' \succ_h \ell$. By Lemma 3(iii), we have that $\mu^{TKDAC}(h) = \ell'$. We conclude that $\mu^{TKDAC}(h) \succ_h \ell$, a contradiction since, by assumption, $\ell \succeq_h \mu^{TKDAC}(h)$. This establishes Claim 3.

We now show that μ^{TKDAC} is weakly envy-free. Towards a contradiction, suppose μ^{TKDAC} is not weakly envy-free. Then, there exists a family $f \in F$ such that ℓ cannot weakly accommodate f alongside $\{g \in F : g \triangleright_\ell f \text{ and } \ell \succeq_g \mu^{TKDAC}(g)\}$. By Claim 3, it follows that ℓ cannot weakly accommodate f alongside $\{g \in F : g \tilde{\triangleright}_\ell f \text{ and } \ell \succeq_g \mu^{TKDAC}(g)\}$. As μ^{TKDAC} is by construction the result of running the TKDA algorithm with the priority profile $\tilde{\triangleright}$, we conclude that the TKDA mechanism is not weakly envy-free, a contradiction of Theorem 6.

$$\mu^{TKDAC} \succeq \mu^{TKDA}$$

We first introduce some additional notation. Let M be the number of rounds of the TKDA algorithm and, for every $i = 1, \dots, M$ and every $\ell \in L$, let Π_ℓ^i be the set of families that propose to ℓ in Round i . Similarly, let \widetilde{M} be the number of rounds of the TKDAC algorithm; that is, the TKDAC algorithm consists of a Clinching Round, which lasts N steps, and then the TKDA algorithm is run with the constructed priority profile $\tilde{\triangleright}$ and lasts \widetilde{M} rounds. For every $i = 1, \dots, \widetilde{M}$ and every $\ell \in L$, let $\widetilde{\Pi}_\ell^i$ be the set of families that propose to ℓ in Round i of the TKDAC algorithm. By definition (Algorithm 5), the threshold of a family f at a locality ℓ only depends on higher-priority families; therefore, $\theta_\ell^f(\Pi_\ell^i \cap \widehat{F}_\ell^f)$ is the threshold of family f for locality ℓ in Round i of the TKDA algorithm and $\theta_\ell^f(\widetilde{\Pi}_\ell^i \cap \widehat{F}_\ell^f(\tilde{\triangleright}))$ is the threshold of family f for locality ℓ in Round i of the TKDAC algorithm.

Claim 4. *In every Round $i = 1, \dots, \widetilde{M}$ of the TKDAC algorithm, every family $f \in F$ proposes to a locality that f weakly prefers to $\mu^{TKDA}(f)$.*

By construction, for every family $f \in F$, $\mu^{TKDAC}(f)$ is the locality to which f proposes in Round \widetilde{M} of the TKDAC algorithm. Therefore, Claim 4 implies that $\mu^{TKDAC}(f) \succeq_f \mu^{TKDA}(f)$ for all $f \in F$, as required. It remains to prove Claim 4.

In Round 1 of the TKDAC algorithm, every family proposes to its most preferred locality; therefore Claim 4 holds for $i = 1$. The remainder of the proof is by induction. We suppose that Claim 4 holds for some $i = 1, \dots, \widetilde{M} - 1$ (induction hypothesis) and show that Claim 4 holds for $i + 1$.

Consider any family $f \in F$ and, for ease of notation, let $\ell = \mu^{TKDA}(f)$. We need to show that, in Round $i + 1$ of the TKDAC algorithm, f proposes to a locality that f weakly prefers to ℓ . By our induction hypothesis, f proposes to a locality that f weakly prefers to ℓ in Round i . If f proposes to a locality that f strictly prefers to ℓ in Round i , then ℓ does not permanently reject f in Round i so f proposes to a locality that f weakly prefers to ℓ in Round $i + 1$, as required. We therefore focus on the case where f proposes to ℓ in Round i of the TKDAC algorithm and need to show that ℓ tentatively accepts f 's proposal. That is, by the construction of the TKDA part of the TKDAC algorithm, we need to show that

$$\theta_\ell^f(\widetilde{\Pi}_\ell^i \cap \widehat{F}_\ell^f) \geq |\widetilde{\Pi}_\ell^i \cap \widehat{F}_\ell^f| + 1. \quad (26)$$

First, as f proposes to ℓ in Round i , $\ell \succeq_f \mu^{TKDAC}(f)$ by construction. Then, by Lemma 3(iii), f does not clinch any locality that f strictly prefers to ℓ in the Clinching Round. By Lemma 2(i), it follows that $\widehat{F}_\ell^f(\triangleright) \subseteq \widehat{F}_\ell^f$.

Second, by construction, family f proposes to and is tentatively accepted by ℓ in the last round of the TKDA algorithm; hence

$$\theta_\ell^f(\Pi_\ell^M \cap \widehat{F}_\ell^f) \geq |\Pi_\ell^M \cap \widehat{F}_\ell^f| + 1. \quad (27)$$

There are two cases: $\theta_\ell^f(\Pi_\ell^M \cap \widehat{F}_\ell^f) = \infty$ and $\theta_\ell^f(\Pi_\ell^M \cap \widehat{F}_\ell^f) \neq \infty$.

Case 1: $\theta_\ell^f(\Pi_\ell^M \cap \widehat{F}_\ell^f) = \infty$. In this case, by the definition of thresholds (Algorithm 5), ℓ can weakly accommodate f alongside \widehat{F}_ℓ^f . As $\widehat{F}_\ell^f(\triangleright) \subseteq \widehat{F}_\ell^f$, ℓ can weakly accommodate f alongside $\widehat{F}_\ell^f(\triangleright)$. Again, by the definition of thresholds (Algorithm 5), we have that $\theta_\ell^f(\widetilde{\Pi}_\ell^i \cap \widehat{F}_\ell^f(\triangleright)) = \infty$; hence inequality (26) is satisfied, as required.

Case 2: $\theta_\ell^f(\Pi_\ell^M \cap \widehat{F}_\ell^f) \neq \infty$. There are two sub-cases: ℓ permanently rejects at least one family with a higher priority than f at ℓ in the TKDA algorithm and ℓ does not permanently reject any family with a higher priority than f at ℓ in the TKDA algorithm.

Sub-case 2.1: ℓ permanently rejects at least one family with a higher priority than f at ℓ in the TKDA algorithm. In this case, let $g \in F$ be the highest-priority family that ℓ permanently rejects in the TKDA algorithm. Therefore, there exists a Round $j = 1, \dots, M$ of the TKDA algorithm, $g \in \Pi_\ell^j$ and

$$\theta_\ell^g(\Pi_\ell^j \cap \widehat{F}_\ell^g) < |\Pi_\ell^j \cap \widehat{F}_\ell^g| + 1. \quad (28)$$

As g is the highest-priority family that ℓ permanently rejects in the TKDA algorithm, any

family with a higher priority that proposes to ℓ in Round j continues to propose to ℓ until the end of the algorithm; therefore, $\Pi_\ell^j \cap \widehat{F}_\ell^g \subseteq \Pi_\ell^M \cap \widehat{F}_\ell^g$.

We next show that the following inequality holds:

$$\theta_\ell^g(\Pi_\ell^M \cap \widehat{F}_\ell^g) < |\Pi_\ell^M \cap \widehat{F}_\ell^g| + 1. \quad (29)$$

First, inequality (29) holds trivially if $\theta_\ell^g(\Pi_\ell^M \cap \widehat{F}_\ell^g) = 0$. Second, inequality (28) implies that $\theta_\ell^g(\Pi_\ell^j \cap \widehat{F}_\ell^g) \neq \infty$; hence Lemma 1(i) implies that $\theta_\ell^g(\Pi_\ell^M \cap \widehat{F}_\ell^g) \neq \infty$. Third, if $\theta_\ell^g(\Pi_\ell^M \cap \widehat{F}_\ell^g) \in \mathbb{Z}_{>0}$, then, as $\Pi_\ell^j \cap \widehat{F}_\ell^g \subseteq \Pi_\ell^M \cap \widehat{F}_\ell^g$, we can apply Lemma 1(ii) to obtain that

$$\theta_\ell^g(\Pi_\ell^M \cap \widehat{F}_\ell^g) \leq \theta_\ell^g(\Pi_\ell^j \cap \widehat{F}_\ell^g) + |\Pi_\ell^M \cap \widehat{F}_\ell^g| - |\Pi_\ell^j \cap \widehat{F}_\ell^g|. \quad (30)$$

Combined with inequality (28), inequality (30) implies inequality (29).

As $g \triangleright_\ell f$, we have that $\Pi_\ell^M \cap \widehat{F}_\ell^g \subseteq \Pi_\ell^M \cap \widehat{F}_\ell^f$. Moreover, as a family's threshold only depends on higher-priority families (Algorithm 5), $\theta_\ell^g(\Pi_\ell^M \cap \widehat{F}_\ell^f) = \theta_\ell^g(\Pi_\ell^M \cap \widehat{F}_\ell^g)$. Combining these observations with inequalities (29) and (27) yields

$$\theta_\ell^g(\Pi_\ell^M \cap \widehat{F}_\ell^f) = \theta_\ell^g(\Pi_\ell^M \cap \widehat{F}_\ell^g) < |\Pi_\ell^M \cap \widehat{F}_\ell^g| + 1 \leq |\Pi_\ell^M \cap \widehat{F}_\ell^f| + 1 \leq \theta_\ell^f(\Pi_\ell^M \cap \widehat{F}_\ell^f). \quad (31)$$

By inequality (28), $\theta_\ell^g(\Pi_\ell^M \cap \widehat{F}_\ell^f) \neq \infty$; hence $\theta_\ell^g(\Pi_\ell^M \cap \widehat{F}_\ell^g) \neq \infty$. Hence, we can apply Lemma 1(iii) to obtain that

$$\theta_\ell^g(\Pi_\ell^M \cap \widehat{F}_\ell^f) \geq \theta_\ell^f(\Pi_\ell^M \cap \widehat{F}_\ell^f),$$

which contradicts inequality (31). We therefore conclude that Sub-case 2.1 cannot occur.

Sub-case 2.2: ℓ does not permanently reject any family with a higher priority than f at ℓ in the TKDA algorithm. In this case, we first show that $\widetilde{\Pi}_\ell^i \cap \widehat{F}_\ell^f \subseteq \Pi_\ell^M \cap \widehat{F}_\ell^f$. Towards a contradiction, suppose that there exists a family $g \in (\widetilde{\Pi}_\ell^i \cap \widehat{F}_\ell^f) \setminus (\Pi_\ell^M \cap \widehat{F}_\ell^f)$. As $(\widetilde{\Pi}_\ell^i \cap \widehat{F}_\ell^f) \setminus (\Pi_\ell^M \cap \widehat{F}_\ell^f) = (\widetilde{\Pi}_\ell^i \setminus \Pi_\ell^M) \cap \widehat{F}_\ell^f$, we have that $g \in \widetilde{\Pi}_\ell^i$, $g \notin \Pi_\ell^M$, and $g \in \widehat{F}_\ell^f$.

First, by our induction hypothesis, g proposes in Round i of the TKDAC algorithm to a locality that g weakly prefers to $\mu^{TKDA}(g)$; therefore, the fact that $g \in \widetilde{\Pi}_\ell^i$ implies that $\ell \succeq_g \mu^{TKDA}(g)$. Second, as the TKDA algorithm matches g to the last locality to which g proposes, the fact that $g \notin \Pi_\ell^M$ implies that $\ell \neq \mu^{TKDA}(g)$. Third, by the assumption of Sub-case 2.2, the fact that $g \in \widehat{F}_\ell^f$ implies that ℓ does not permanently reject g in the TKDA algorithm, hence $\mu^{TKDA}(g) \succeq_g \ell$. Combining these observations, we have that $\ell \succeq_g \mu^{TKDA}(g)$, $\ell \neq \mu^{TKDA}(g)$, and $\mu^{TKDA}(g) \succeq_g \ell$, a contradiction.

Having established that $\widetilde{\Pi}_\ell^i \cap \widehat{F}_\ell^f \subseteq \Pi_\ell^M \cap \widehat{F}_\ell^f$ and that $\widehat{F}_\ell^f(\bowtie) \subseteq \widehat{F}_\ell^f$, we conclude that $\widetilde{\Pi}_\ell^i \cap \widehat{F}_\ell^f(\bowtie) \subseteq \Pi_\ell^M \cap \widehat{F}_\ell^f$. It follows that $|\widetilde{\Pi}_\ell^i \cap \widehat{F}_\ell^f(\bowtie)| \leq |\Pi_\ell^M \cap \widehat{F}_\ell^f|$. Moreover, as $\theta_\ell^f(\Pi_\ell^M \cap \widehat{F}_\ell^f) \neq \infty$ by the assumption of Case 2, and as $\theta_\ell^f(\Pi_\ell^M \cap \widehat{F}_\ell^f) \neq 0$ by inequality (27), we have that $\theta_\ell^f(\Pi_\ell^M \cap \widehat{F}_\ell^f) \in \mathbb{Z}_{>0}$ so we can apply Lemma 1(ii) to obtain that

$$\theta_\ell^f(\Pi_\ell^M \cap \widehat{F}_\ell^f) \leq \theta_\ell^f(\widetilde{\Pi}_\ell^i \cap \widehat{F}_\ell^f(\bowtie)) + |\Pi_\ell^M \cap \widehat{F}_\ell^f| - |\widetilde{\Pi}_\ell^i \cap \widehat{F}_\ell^f(\bowtie)|. \quad (32)$$

Combining inequality (32) with inequality (27) yields inequality (26), as required. \square

Proof of Lemma 2

We prove the lemma by a single induction argument. To show that the lemma holds for $j = 1$, note that (i) $\tilde{\triangleright}^0 = \tilde{\triangleright}^1 = \triangleright$, (ii) $\emptyset = \Delta_f^0 \subseteq \Delta_f^1$ for every $f \in F$, (iii) $\emptyset = \Gamma_f^0 \subseteq \Gamma_f^1$ for every $f \in F$, and (iv) $\emptyset = \Theta_\ell^0 \subseteq \Theta_\ell^1$ for every $\ell \in L$.

For the induction step, let us assume that part (iv) of the lemma holds for some $j = 1, \dots, N$, i.e., $\Theta_\ell^{j-1} \subseteq \Theta_\ell^j$ for every $\ell \in L$ and every $j = 1, \dots, N$. We will show that the assumption implies parts (i)-(iv) of the lemma in Step $j+1$. That is, we consider an arbitrary family $f \in F$ and an arbitrary locality $\ell \in L$ and show the following:

- (i) If $f \notin \Theta_{\ell'}^j$ for all $\ell' \in L$ such that $\ell' \succ_f \ell$, then $\widehat{F}_\ell^f(\tilde{\triangleright}^{j+1}) \subseteq \widehat{F}_\ell^f(\tilde{\triangleright}^j)$;
- (ii) $\Delta_f^j \subseteq \Delta_f^{j+1}$;
- (iii) If $\ell \in \Gamma_f^j \setminus \Gamma_f^{j+1}$, then there exists $\ell' \in L$ such that $\ell' \succ_f \ell$ and $f \in \Theta_{\ell'}^j$; and
- (iv) $\Theta_\ell^j \subseteq \Theta_\ell^{j+1}$.

Proof of (i) We prove the contrapositive. Suppose there exists a family $g \in \widehat{F}_\ell^f(\tilde{\triangleright}^{j+1}) \setminus \widehat{F}_\ell^f(\tilde{\triangleright}^j)$, then $f \tilde{\triangleright}_\ell^j g$ and $g \tilde{\triangleright}_\ell^{j+1} f$. We need to show that f clinches a locality that f strictly prefers to ℓ in Step j . There are two cases: $f \triangleright_\ell g$ and $g \triangleright_\ell f$.

Case 1: $f \triangleright_\ell g$. Since $g \tilde{\triangleright}_\ell^{j+1} f$, then by construction (Step j (e) of the Clinching Round) we have that f clinches a locality that f strictly prefers to ℓ in Step j (c) of the Clinching Round.

Case 2: $g \triangleright_\ell f$. Since $f \tilde{\triangleright}_\ell^j g$, then by construction (Step j (e) of the Clinching Round), g clinches a locality that g strictly prefers to ℓ in Step $j-1$ (c) of the Clinching Round. Since we have assumed that $\Theta_\ell^{j-1} \subseteq \Theta_\ell^j$, g continues to clinch that locality in Step j (c) of the Clinching Round. Since $g \tilde{\triangleright}_\ell^{j+1} f$, we have that f clinches a locality that f strictly prefers to ℓ in Step j (c) of the Clinching Round.

Proof of (ii) Suppose that $\ell \in \Delta_f^j$, we need to show that $\ell \in \Delta_f^{j+1}$. As $\ell \in \Delta_f^j$, ℓ cannot weakly accommodate f alongside $\Theta_\ell^{j-1} \cap \widehat{F}_\ell^f(\tilde{\triangleright}^j)$. Similarly, $\ell \in \Delta_f^{j+1}$ if ℓ cannot weakly accommodate f alongside $\Theta_\ell^j \cap \widehat{F}_\ell^f(\tilde{\triangleright}^{j+1})$. Therefore, it is sufficient to show that

$$(\Theta_\ell^{j-1} \cap \widehat{F}_\ell^f(\tilde{\triangleright}^j)) \subseteq (\Theta_\ell^j \cap \widehat{F}_\ell^f(\tilde{\triangleright}^{j+1})).$$

Towards a contradiction, suppose that there exists a family

$$g \in (\Theta_\ell^{j-1} \cap \widehat{F}_\ell^f(\tilde{\triangleright}^j)) \setminus (\Theta_\ell^j \cap \widehat{F}_\ell^f(\tilde{\triangleright}^{j+1})).$$

Since we have assumed that $\Theta_\ell^{j-1} \subseteq \Theta_\ell^j$, we have that

$$g \in (\Theta_\ell^{j-1} \cap \widehat{F}_\ell^f(\tilde{\triangleright}^j)) \setminus \widehat{F}_\ell^f(\tilde{\triangleright}^{j+1}) = \Theta_\ell^{j-1} \cap (\widehat{F}_\ell^f(\tilde{\triangleright}^j) \setminus \widehat{F}_\ell^f(\tilde{\triangleright}^{j+1})).$$

Since $g \in \widehat{F}_\ell^f(\tilde{\triangleright}^j) \setminus \widehat{F}_\ell^f(\tilde{\triangleright}^{j+1})$, we have that $g \tilde{\triangleright}_\ell^j f$ and $f \tilde{\triangleright}_\ell^{j+1} g$. Therefore, $f \in \widehat{F}_\ell^g(\tilde{\triangleright}_\ell^{j+1}) \setminus \widehat{F}_\ell^g(\tilde{\triangleright}^j)$ and, as a result, $\widehat{F}_\ell^g(\tilde{\triangleright}_\ell^{j+1}) \not\subseteq \widehat{F}_\ell^g(\tilde{\triangleright}^j)$. Then, the contrapositive of (i) implies that $g \in \Theta_{\ell'}^j$ for

some $\ell' \in L$ such that $\ell' \succ_g \ell$. Since $g \in \Theta_\ell^{j-1}$ and we have assumed that $\Theta_\ell^{j-1} \subseteq \Theta_\ell^j$, it follows that g clinches both ℓ and ℓ' in Step j . This is a contradiction since, by construction (Step j (c) of the Clinching Round), a family can clinch at most one locality in any given step of the Clinching Round.

Proof of (iii) Suppose that $\ell \in \Gamma_f^j \setminus \Gamma_f^{j+1}$. Then, ℓ can weakly accommodate f alongside $\widehat{F}_\ell^f(\widetilde{\mathcal{L}}^j)$ but not alongside $\widehat{F}_\ell^f(\widetilde{\mathcal{L}}^{j+1})$, which implies that $\widehat{F}_\ell^f(\widetilde{\mathcal{L}}^{j+1}) \not\subseteq \widehat{F}_\ell^f(\widetilde{\mathcal{L}}^j)$. By the contrapositive of (i), we have that $f \in \Theta_{\ell'}^j$ for some $\ell' \in L$ such that $\ell' \succ_f \ell$.

Proof of (iv) Suppose that $f \in \Theta_\ell^j$, we need to show that $f \in \Theta_\ell^{j+1}$. By construction (Step j (c) of Clinching Round), ℓ proposes to f in Step j (b) and all localities that f prefers to ℓ reject f in Step j (a); that is, $\ell \in \Gamma_f^j$ and, for all $\ell' \in L$ such that $\ell' \succ_f \ell$, $\ell' \in \Delta_f^j$. Since a family can clinch at most one locality in each step, $f \notin \Theta_{\ell'}^j$ for all $\ell' \in L$ such that $\ell' \succ_f \ell$. By the contrapositive of (iii), $\ell \notin \Gamma_f^j \setminus \Gamma_f^{j+1}$; therefore $\ell \in \Gamma_f^{j+1}$. Moreover, by (ii), $\ell' \in \Delta_f^{j+1}$ for all $\ell' \in L$ such that $\ell' \succ_f \ell$. Therefore, ℓ proposes to f in Step $j+1$ and all localities that f prefers to ℓ reject f in Step $j+1$, which means that f clinches ℓ in Step $j+1$, i.e., $f \in \Theta_\ell^{j+1}$. \square

Proof of Lemma 3

For ease of notation, let $\mu = \mu^{TKDAC}$. We consider a Step $j = 1, \dots, N$ of the Clinching Round, a family $f \in F$, and a locality $\ell \in L$. We prove each part of the lemma in turn.

Proof of (i) We first show that if a family g is rejected by $\mu(g)$ in Step $i = 2, \dots, N$ of the Clinching Round, then a family h is rejected by $\mu(h)$ in Step $i-1$.

Suppose that a family $g \in F$ is rejected by $\mu(g)$ in Step $i = 2, \dots, N$ of the Clinching Round. Since $\mu(g)$ rejects g in Step i (a) of the Clinching Round, $\mu(g)$ cannot weakly accommodate g alongside $\Theta_{\mu(g)}^{i-1} \cap \widehat{F}_{\mu(g)}^g(\widetilde{\mathcal{L}}^i)$. If all families in $\Theta_{\mu(g)}^{i-1} \cap \widehat{F}_{\mu(g)}^g(\widetilde{\mathcal{L}}^i)$ are matched to $\mu(g)$ at the end of the TKDAC algorithm (i.e., if $(\Theta_{\mu(g)}^{i-1} \cap \widehat{F}_{\mu(g)}^g(\widetilde{\mathcal{L}}^i)) \subseteq \mu(\mu(g))$), then $\mu(g)$ can accommodate $(\Theta_{\mu(g)}^{i-1} \cap \widehat{F}_{\mu(g)}^g(\widetilde{\mathcal{L}}^i)) \cup \{g\}$. Therefore, $\mu(g)$ can weakly accommodate g alongside $\Theta_{\mu(g)}^{i-1} \cap \widehat{F}_{\mu(g)}^g(\widetilde{\mathcal{L}}^i)$, a contradiction. We conclude that there exists a family $h \in \Theta_{\mu(g)}^{i-1} \cap \widehat{F}_{\mu(g)}^g(\widetilde{\mathcal{L}}^i)$ such that $\mu(h) \neq \mu(g)$. Since $h \in \Theta_{\mu(g)}^{i-1}$, we have that $h \in \Theta_{\mu(g)}^N$ by Lemma 2(iv). It follows that $\mu(g)$ proposes to h in Step N of the Clinching Round, which in turn implies that $\mu(g)$ can weakly accommodate h alongside $\widehat{F}_\ell^h(\widetilde{\mathcal{L}}^N) = \widehat{F}_\ell^h(\widetilde{\mathcal{L}})$. Then, in every round of the TKDAC algorithm, h 's threshold at $\mu(g)$ is ∞ (see Algorithm 5). It follows that $\mu(g)$ tentatively accepts any proposal from h so the fact that $\mu(h) \neq \mu(g)$ implies that h does not propose to $\mu(g)$ in the TKDAC algorithm, hence $\mu(h) \succ_h \mu(g)$. Finally, since $h \in \Theta_{\mu(g)}^{i-1}$ and $\mu(h) \succ_f \mu(g)$, $\mu(h)$ rejects h in Step $i-1$ (c) of the Clinching Round, as required.

Now, suppose towards a contradiction that $\mu(f) \in \Delta_f^j$. The preceding argument implies, by induction, that there exists a family f' such that $\mu(f') \in \Delta_{f'}^1$, i.e., $\mu(f')$ rejects f' in Step 1 of the Clinching Round. Then, $\mu(f')$ cannot weakly accommodate f' alongside $\Theta_{\mu(f')}^0 \cap \widehat{F}_{\mu(f')}^{f'}(\widetilde{\mathcal{L}}^1)$. Since $\Theta_{\mu(f')}^0 = \emptyset$, it follows that $\mu(f')$ cannot weakly accommodate f' on its own. Therefore, the threshold of f' at $\mu(f')$ is 0 in every round of the TKDAC algorithm.

We conclude that $\mu(f')$ permanently rejects any proposal from f' , which contradicts the assumption that f' is matched to $\mu(f')$ at the end of the TKDAC algorithm.

Proof of (ii) Suppose that $\ell \in \Gamma_f^j$. We need to show that $\mu(f) \succeq_f \ell$. Then, by Lemma 2(iii), $\ell \in \Gamma_f^j$ implies that one of the following two cases holds: either $\ell \in \Gamma_f^N$ or $f \in \Theta_{\ell'}^{j-1}$ for some $\ell' \in L$ such that $\ell' \succ_f \ell$.

Case 1: $\ell \in \Gamma_f^N$. In this case, by construction (Step $N(c)$ of the Clinching Round), ℓ can weakly accommodate f alongside $\widehat{F}_\ell^f(\widehat{\succ}^N) = \widehat{F}_\ell^f(\widehat{\succ})$. Therefore, in every round of the TKDAC algorithm, f 's threshold at ℓ is ∞ (see Algorithm 5). It follows that ℓ tentatively accepts any proposal from f in the TKDAC algorithm, hence $\mu(f) \succeq_f \ell$, as required.

Case 2: $f \in \Theta_{\ell'}^{j-1}$ for some $\ell' \in L$ such that $\ell' \succ_f \ell$. In this case, Lemma 2(iv), we have that $f \in \Theta_{\ell'}^N$ for some $\ell' \in L$ such that $\ell' \succ_f \ell$. Then, by construction (Step $N(c)$ of the Clinching Round), ℓ' proposes to f in Step N of the Clinching Round, i.e., $\ell' \in \Gamma_f^N$. Then, f 's threshold at ℓ' is ∞ in every round of the TKDAC algorithm, hence $\mu(f) \succeq_f \ell' \succ_f \ell$, as required.

Proof of (iii) Suppose that $f \in \Theta_\ell^j$. Then, by construction (Step $j(c)$ of the Clinching Round) we have that $\ell \in \Gamma_f^j$ and, for all $\ell' \in L$ such that $\ell' \succ_f \ell$, it is the case that $\ell' \in \Delta_f^j$. From parts (i) and (ii) of the lemma, we obtain that $\mu(f) \succeq_f \ell$ and that $\mu(f) \neq \ell'$ for all $\ell' \in L$ such that $\ell' \succ_f \ell$. We conclude that $\mu(f) = \ell$, as required. \square

Proof of Lemma 4

We prove the first part of the lemma by a single induction argument. For the initial step, we have that $\Theta_\ell^0 = \overline{\Theta}_\ell^0 = \emptyset$ for all $\ell \in L$.

For the induction step, let us assume that $\Theta_\ell^{j-1} \setminus \{f\} = \overline{\Theta}_\ell^{j-1} \setminus \{f\}$ holds for some $j = 1, \dots, q$ and all $\ell \in L$. We will show that $\widetilde{\succ}^j = \overline{\succ}^j$, $\widetilde{\Delta}_g^j = \overline{\Delta}_g^j$, $\widetilde{\Gamma}_g^j = \overline{\Gamma}_g^j$, and, finally, that $\Theta_\ell^j \setminus \{f\} = \overline{\Theta}_\ell^j \setminus \{f\}$ for all $g \in F$ and $\ell \in L$.

Since $j \leq q$, we have that $j - 1 < \min\{m, \overline{m}\}$. Therefore, f does not clinch any locality in Step $j - 1$ of the Clinching Round with either report; therefore, $\Theta_\ell^{j-1} \setminus \{f\} = \overline{\Theta}_\ell^{j-1} \setminus \{f\}$ for all $\ell \in L$ implies that $\Theta_\ell^{j-1} = \overline{\Theta}_\ell^{j-1}$ for all $\ell \in L$. In Step $j - 1(e)$ of the Clinching Round the construction of $\widetilde{\succ}^j$ only depends on which families have clinched which localities, hence $\Theta_\ell^{j-1} = \overline{\Theta}_\ell^{j-1}$ for all $\ell \in L$ implies that $\widetilde{\succ}^j = \overline{\succ}^j$. In Step $j(a)$, the fact that $\Theta_\ell^{j-1} = \overline{\Theta}_\ell^{j-1}$ for all $\ell \in L$ and $\widetilde{\succ}^j = \overline{\succ}^j$ implies that every locality rejects the same families under both reports so $\Delta_g^j = \overline{\Delta}_g^j$ for all $g \in F$. Similarly, in Step $j(b)$, $\widetilde{\succ}^j = \overline{\succ}^j$ implies that $\Gamma_g^j = \overline{\Gamma}_g^j$ for all $g \in F$. Finally, consider any family $h \neq f$ in Step $j(c)$. As $\Delta_h^j = \overline{\Delta}_h^j$, $\Gamma_h^j = \overline{\Gamma}_h^j$, and h does not misreport its preferences (only f 's report changes from \succ_f to \succ'_f), h clinches the same locality (if any), whether f reports \succ_f or \succ'_f . We therefore conclude that $\Theta_\ell^j \setminus \{f\} = \overline{\Theta}_\ell^j \setminus \{f\}$ for all $\ell \in L$, as required.

We now turn to the second part of the lemma. Consider any $j = 1, \dots, q$ with $j < \{m, \overline{m}\}$. We have established that $\Theta_\ell^j \setminus \{f\} = \overline{\Theta}_\ell^j \setminus \{f\}$ for all $\ell \in L$. As $j < \{m, \overline{m}\}$, whether f

reports \succ_f or \succ'_f , f does not clinch any locality in Step j of the Clinching Round. Therefore $\Theta_\ell^j = \bar{\Theta}_\ell^j$ for all $\ell \in L$, as required. \square

B Additional examples

B.1 Non-existence of stable outcomes

We reproduce an example from [McDermid and Manlove \(2010\)](#). There are three families, two localities, and two dimensions. The family sizes and locality capacities are

$$\nu = d_1 \begin{pmatrix} f_1 & f_2 & f_3 \\ 1 & 1 & 2 \end{pmatrix} \quad \text{and} \quad \kappa = d_1 \begin{pmatrix} \ell_1 & \ell_2 \\ 2 & 1 \end{pmatrix}.$$

The preferences and priorities are:

$$\succ_{f_1}: \ell_2, \ell_1, \emptyset \quad \succ_{f_2}: \ell_1, \ell_2, \emptyset \quad \succ_{f_3}: \ell_1, \emptyset, \ell_2 \quad \triangleright_{\ell_1}: f_1, f_3, f_2 \quad \triangleright_{\ell_2}: f_2, f_1, f_3.$$

Suppose, towards a contradiction, that there exists a stable matching μ in this example. Since ℓ_2 cannot accommodate f_3 , either $\mu(f_3) = \ell_1$ or $\mu(f_3) = \emptyset$. If $\mu(f_3) = \ell_1$, then $\mu(f_2) = \ell_2$ as otherwise f_2 and ℓ_2 form a blocking pair. Then, $\mu(f_1) = \emptyset$ so f_1 and ℓ_1 form a blocking pair and μ is not stable. If $\mu(f_3) = \emptyset$, then $\mu(f_1) = \ell_1$ as otherwise f_1 and ℓ_1 form a blocking pair. In turn, f_2 and ℓ_1 form a blocking pair unless $\mu(f_2) = \ell_1$. However, this matching is not stable since f_2 and ℓ_2 form a blocking pair.

One way to understand this negative result is that the choice function induced by ℓ_1 's priorities does not satisfy substitutability ([Roth, 1984b](#)) because of a complementarity between f_1 and f_2 . If only f_2 and f_3 compete for ℓ_1 , stability dictates that f_3 be accepted and f_2 rejected. If f_1 is also competing, stability dictates that f_1 and f_2 be accepted and f_3 rejected. Thus, there is a complementarity between f_1 and f_2 in the sense that f_2 is accepted when f_1 is also considered but not otherwise.

B.2 Knapsack Top Trading Cycles with Endowment

We illustrate the KTTCE algorithm using the example from the proof of [Theorem 1](#). We add the following lexicographic priorities:

$$\triangleright_{\ell_1}: f_1, f_2, \dots \quad \triangleright_{\ell_2}: \mathbf{f_1}, \mathbf{f_2}, \dots \quad \triangleright_{\ell_3}: \mathbf{f_3}, \dots \quad \triangleright_{\ell_4}: \mathbf{f_4}, \dots$$

The KTTCE algorithm also depends on the order in which families are picked, should the algorithm enter the Rejection Stage. With four families, there are $4! = 24$ such orderings; however, we will see that the only part of that ordering that matters for the outcome is whether f_3 or f_4 is picked first. Therefore, we describe the KTTCE algorithm under two different orderings: KTTCE3 picks f_3 before f_4 and KTTCE4 picks f_4 before f_3 .

KTTCE3

The workings of KTTCE3 are displayed in [Figure 8](#). In Round 1, the unique trading cycle is $f_1 \rightarrow \ell_3 \rightarrow f_3 \rightarrow \ell_2 \rightarrow f_1$. The unique trading cycle is not feasible since ℓ_2 cannot accommodate f_3 alongside f_2 . Therefore, the algorithm enters the Rejection Stage. As f_1

and f_2 require only one unit of capacity, they can replace any family at any locality; hence no permanent rejection occurs if one of these families is picked. In contrast, neither f_3 nor f_4 can replace f_1 at ℓ_2 . By assumption, f_3 is picked before f_4 and permanently rejected by ℓ_2 .

In Round 2, f_3 points at its second preference ℓ_3 and the feasible (and trivial) cycle $f_3 \rightarrow \ell_3 \rightarrow f_3$ appears so f_3 is permanently matched to ℓ_3 . As a result, ℓ_3 is “full” and permanently rejects all other families, including f_1 . In Round 3, f_1 points at its second preference ℓ_1 and is permanently matched to it, effectively taking advantage of ℓ_1 ’s unassigned unit of capacity. In Round 4, since f_1 has been permanently matched, ℓ_2 points at its second-priority family f_2 . In addition, $\mu^A(\ell_2) = \emptyset$ since f_1 has left ℓ_2 to be permanently matched to ℓ_1 . As a result, the cycle $f_2 \rightarrow \ell_4 \rightarrow f_4 \rightarrow \ell_2 \rightarrow f_2$ is feasible so both families are permanently matched to the locality at which they are pointing. Since all families are permanently matched, the algorithm ends and produces the following matching:

$$\begin{pmatrix} f_1 & f_2 & f_3 & f_4 \\ \ell_1 & \ell_4 & \ell_3 & \ell_2 \end{pmatrix}.$$

KTTCE4

The workings of KTTCE4 are displayed in Figure 9. As for KTTCE3, the unique cycle in Round 1 is $f_1 \rightarrow \ell_3 \rightarrow f_3 \rightarrow \ell_2 \rightarrow f_1$, which is not feasible. The difference in KTTCE4 compared to KTTCE3 is that in the Rejection Stage f_4 is picked and permanently rejected by ℓ_2 . In Round 2, f_4 points at its second preference ℓ_4 (which points back at f_4) so f_4 is permanently matched to ℓ_4 . In Round 3, the unique cycle is again $f_1 \rightarrow \ell_3 \rightarrow f_3 \rightarrow \ell_2 \rightarrow f_1$, which is still infeasible. This time, ℓ_2 permanently rejects f_3 (because f_3 must be picked in the Rejection Stage). In Round 4, f_3 points at and is permanently matched to ℓ_3 . In Round 5, ℓ_3 is “full” and permanently rejects f_1 , which points at and is permanently matched to ℓ_1 . In Round 6 (not displayed), f_2 and ℓ_2 points at one another and are permanently matched. The algorithm ends and produces the following matching:

$$\begin{pmatrix} f_1 & f_2 & f_3 & f_4 \\ \ell_1 & \ell_2 & \ell_3 & \ell_4 \end{pmatrix}.$$

Discussion

Observe first that both matchings Pareto dominate the endowment. This is not surprising since there is only one dimension and the priorities are lexicographic, hence Theorem 4 applies. Perhaps more surprising is the fact that KTTCE3 produces a chain-efficient matching, which may appear at odds with Theorem 1. What is more, we showed in the proof of Theorem 1 that no individually rational, chain-efficient, and strategy-proof mechanism exists in this specific market. However, the fact that KTTCE3 produces a chain-efficient matching in this instance does not mean it is a chain-efficient mechanism. In fact, one can show that, if f_1 reports its preferences to be ℓ_3, ℓ_2, \dots , KTTCE3 produces

$$\begin{pmatrix} f_1 & f_2 & f_3 & f_4 \\ \ell_2 & \ell_1 & \ell_3 & \ell_4 \end{pmatrix},$$

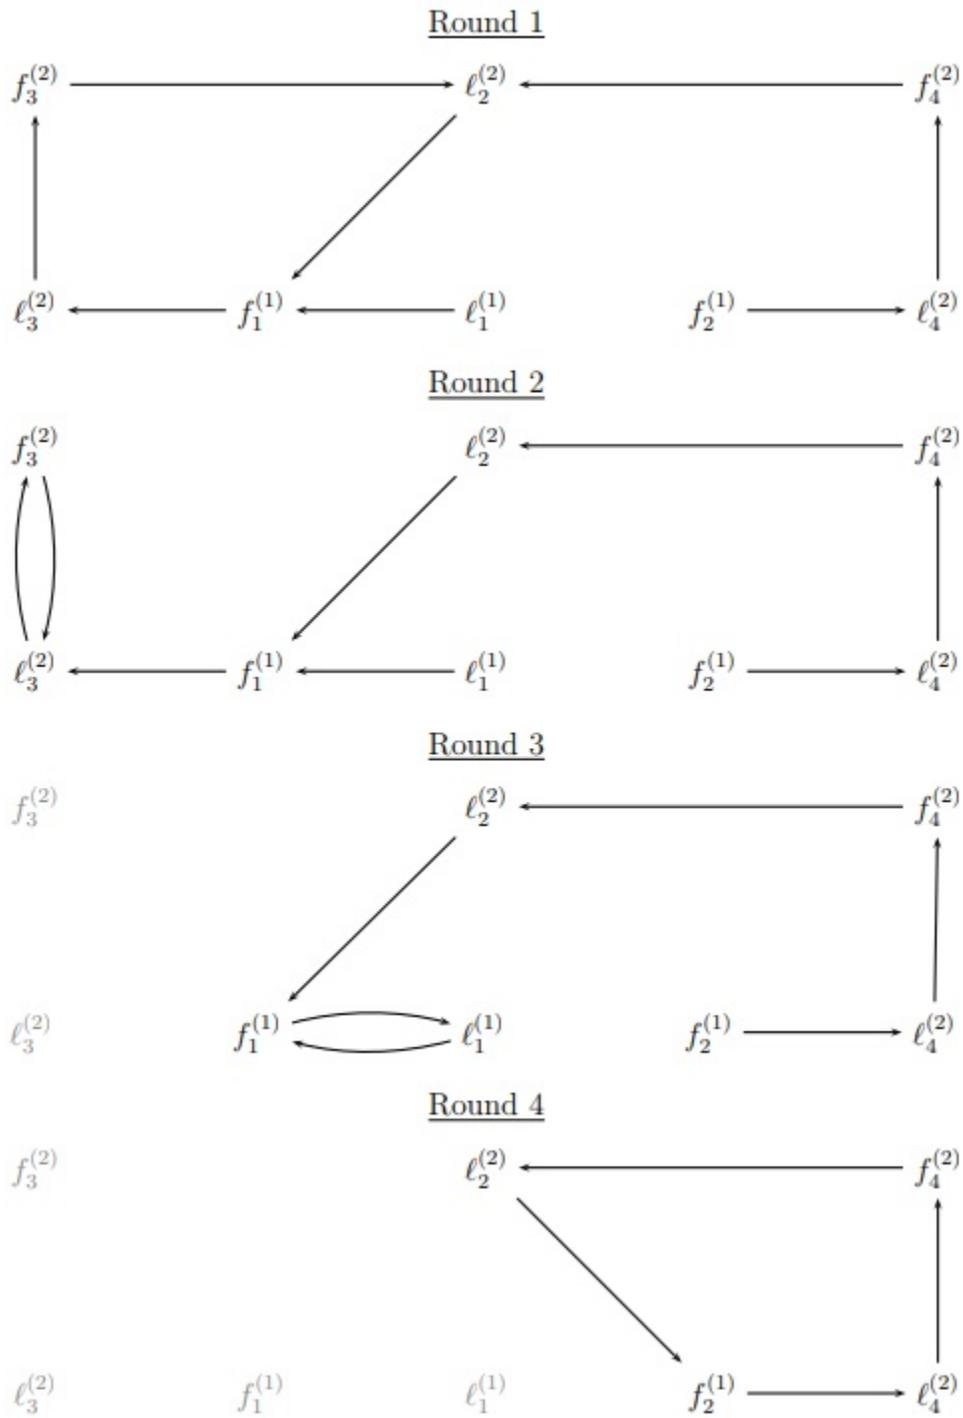


Figure 8: Workings of KTTCE3.

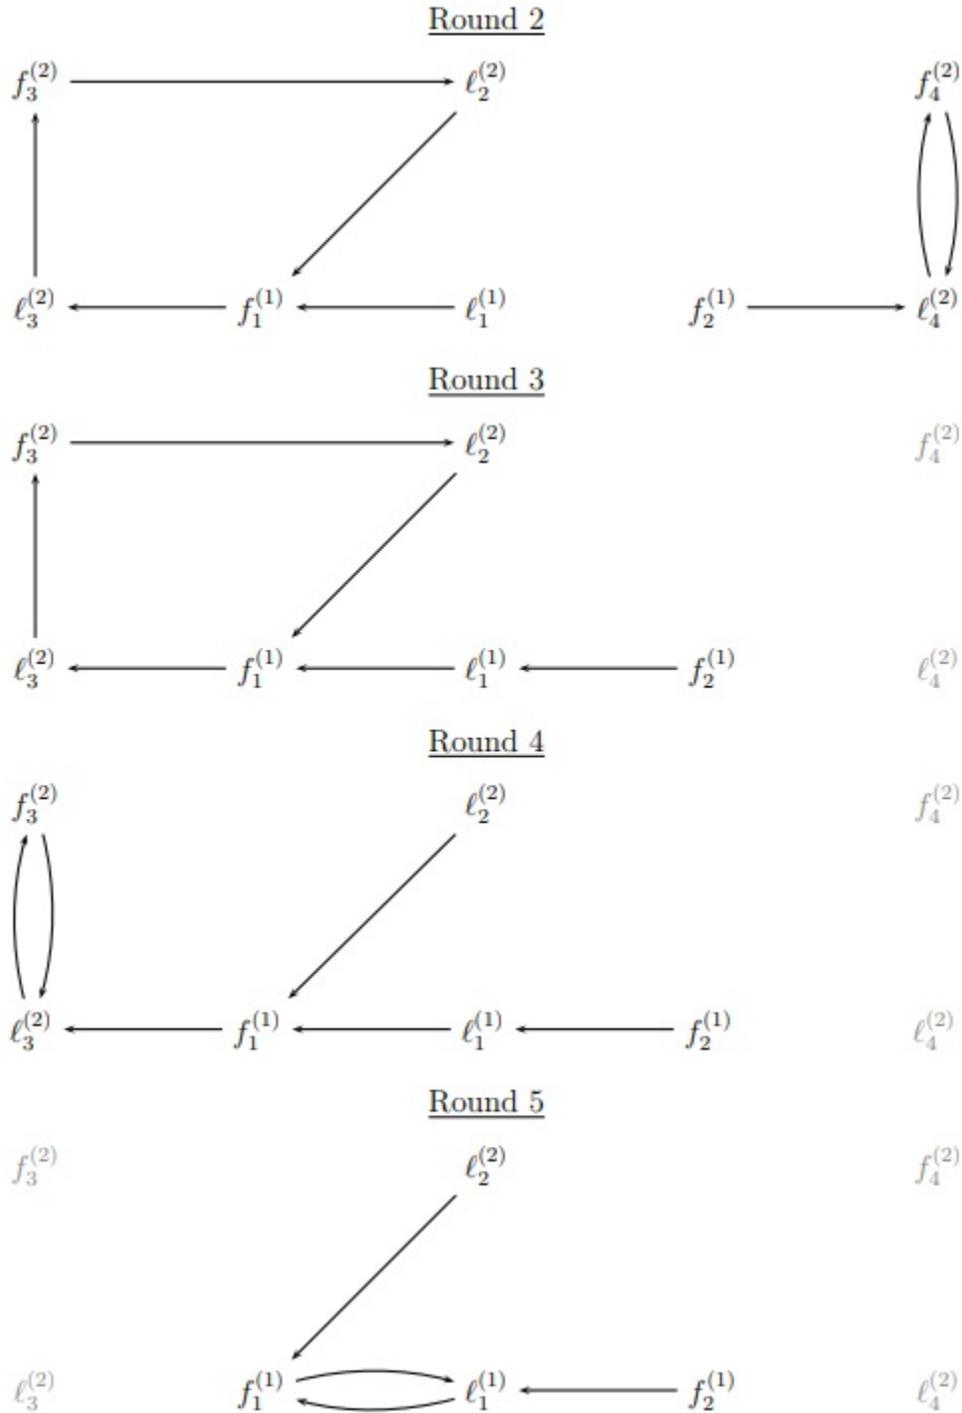


Figure 9: Workings of KTTCE4.

which is not chain-efficient.

Second, the matching produced by KTTCE3 Pareto dominates the one produced by KTTCE4. This is due to the fact that picking f_3 allows the algorithm to match f_1 to ℓ_1 in Round 3, which makes the cycle $f_2 \rightarrow \ell_4 \rightarrow f_4 \rightarrow \ell_2 \rightarrow f_2$ feasible in Round 4. In contrast, picking f_4 in KTTCE4 does not allow matching f_2 to ℓ_1 in Round 3 (since ℓ_1 points at f_1), and therefore the cycle $f_1 \rightarrow \ell_3 \rightarrow f_3 \rightarrow \ell_2 \rightarrow f_1$ remains infeasible. One might therefore wonder whether the picking order can be designed in a way that maximizes the efficiency of the mechanism. This would require picking the “best family” (from an efficiency point of view) every time the algorithm enters the Rejection Stage. Unfortunately, what constitutes the best family depends on preferences; therefore such a mechanism would violate strategy-proofness. In order for the mechanism to be strategy-proof, the picking order must be entirely independent of preferences, which has an efficiency cost. Third, the KTTCE mechanism may produce different outcomes with different priorities. If the priorities are $\succ_{\ell_1}: f_1, f_2, \dots$ and $\succ_{\ell_2}: f_2, f_1, \dots$ (without changing the priorities of ℓ_3 and ℓ_4 so that priorities remain lexicographic), KTTCE3 and KTTCE4 produce the same matchings as they do in our example above. If the priorities are either $\succ_{\ell_1}: f_2, f_1, \dots$ and $\succ_{\ell_2}: f_2, f_1, \dots$ or $\succ_{\ell_1}: f_2, f_1, \dots$ and $\succ_{\ell_2}: f_1, f_2, \dots$, KTTCE3 and KTTCE4 respectively produce the following matchings:

$$\begin{pmatrix} f_1 & f_2 & f_3 & f_4 \\ \ell_2 & \ell_1 & \ell_3 & \ell_4 \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} f_1 & f_2 & f_3 & f_4 \\ \ell_3 & \ell_1 & \ell_2 & \ell_4 \end{pmatrix}.$$

B.3 Threshold Knapsack Deferred Acceptance

We illustrate the TKDA algorithm with the following example. There are seven families, four localities, and two dimensions.

Preferences:

$$\begin{aligned} \succ_{f_1}: \ell_2, \ell_3, \dots & \quad \succ_{f_2}: \ell_4, \ell_1, \dots & \quad \succ_{f_3}: \ell_2, \ell_1, \dots & \quad \succ_{f_4}: \ell_1, \dots \\ \succ_{f_5}: \ell_1, \ell_2, \dots & \quad \succ_{f_6}: \ell_1, \ell_2, \ell_3, \dots & \quad \succ_{f_7}: \ell_4, \dots \end{aligned}$$

Priorities:

$$\begin{aligned} \triangleright_{\ell_1}: f_1, f_2, f_3, f_4, f_5, f_6, f_7 & \quad \triangleright_{\ell_2}: f_5, f_1, f_2, f_7, f_6, f_3, f_4 \\ \triangleright_{\ell_3}: f_4, f_6, f_1, \dots & \quad \triangleright_{\ell_4}: f_7, f_2, \dots \end{aligned}$$

Sizes and capacities:

$$\nu = \begin{matrix} & f_1 & f_2 & f_3 & f_4 & f_5 & f_6 & f_7 \\ \begin{matrix} d_1 \\ d_2 \end{matrix} & \begin{pmatrix} 2 & 1 & 1 & 2 & 1 & 1 & 1 \\ 1 & 0 & 0 & 0 & 1 & 0 & 0 \end{pmatrix} \end{matrix} \quad \kappa = \begin{matrix} & \ell_1 & \ell_2 & \ell_3 & \ell_4 \\ \begin{matrix} d_1 \\ d_2 \end{matrix} & \begin{pmatrix} 4 & 4 & 3 & 1 \\ 2 & 1 & 1 & 0 \end{pmatrix} \end{matrix}.$$

We show that, in this example, the TKDA algorithm produces the following matching:

$$\mu^{TKDA} = \begin{pmatrix} f_1 & f_2 & f_3 & f_4 & f_5 & f_6 & f_7 \\ \ell_3 & \ell_1 & \ell_1 & \ell_1 & \ell_2 & \ell_3 & \ell_4 \end{pmatrix}.$$

The TKDA algorithm lasts four rounds. The first three rounds are displayed in Table 10. In Round 1, every family proposes to its first-preference locality. Families f_4 , f_5 , and f_6

TKDA - Round 1

l_1	(4, 2)	θ	l_2	(4, 1)	θ	l_3	(3, 1)	θ	l_4	(1, 0)	θ
f_1	(2, 1)	∞	f_5	(1, 1)	∞	f_4	(2, 0)	∞	f_7	(1, 0)	∞
f_2	(1, 0)	∞	f_1	(2, 1)	1	f_6	(1, 0)	∞	f_8	(1, 0)	0
f_3	(1, 0)	∞	f_2	(1, 0)	∞	f_1	(2, 1)	1	\vdots		
f_4	(2, 0)	2	f_7	(1, 0)	1	\vdots					
f_5	(1, 1)	2	f_6	(1, 0)	1						
f_8	(1, 0)	2	f_8	(1, 0)	1						
f_7	(1, 0)	0	f_4	(2, 0)	0						

TKDA - Round 2

l_1	(4, 2)	θ	l_2	(4, 1)	θ	l_3	(3, 1)	θ	l_4	(1, 0)	θ
f_1	(2, 1)	∞	f_5	(1, 1)	∞	f_4	(2, 0)	∞	f_7	(1, 0)	∞
f_2	(1, 0)	∞	f_1	(2, 1)	1	f_6	(1, 0)	∞	f_2	(1, 0)	0
f_3	(1, 0)	∞	f_2	(1, 0)	∞	f_1	(2, 1)	1	\vdots		
f_4	(2, 0)	3	f_7	(1, 0)	1	\vdots					
f_8	(1, 1)	0	f_8	(1, 0)	1						
f_6	(1, 0)	0	f_3	(1, 0)	1						
f_7	(1, 0)	0	f_4	(2, 0)	0						

TKDA - Round 3

l_1	(4, 2)	θ	l_2	(4, 1)	θ	l_3	(3, 1)	θ	l_4	(1, 0)	θ
f_1	(2, 1)	∞	f_5	(1, 1)	∞	f_4	(2, 0)	∞	f_7	(1, 0)	∞
f_2	(1, 0)	∞	f_8	(2, 1)	0	f_6	(1, 0)	∞	f_2	(1, 0)	0
f_3	(1, 0)	∞	f_2	(1, 0)	∞	f_1	(2, 1)	2	\vdots		
f_4	(2, 0)	3	f_7	(1, 0)	0	\vdots					
f_5	(1, 1)	0	f_6	(1, 0)	0						
f_6	(1, 0)	0	f_3	(1, 0)	0						
f_7	(1, 0)	0	f_4	(2, 0)	0						

Table 10: Rounds 1-3 of the TKDA algorithm. Sizes and capacities in parentheses. f_i : f_i proposes and is tentatively accepted. ~~f_i~~ : f_i proposes and is permanently rejected.

propose to ℓ_1 . The threshold at ℓ_1 of f_1 , f_2 , and f_3 is ∞ since ℓ_1 can accommodate all three families together. Family f_4 's threshold at ℓ_1 is 2 because ℓ_1 can weakly accommodate f_4 alongside any one of f_1 , f_2 , or f_3 , but not alongside $\{f_1, f_2\}$ or $\{f_1, f_3\}$. Family f_5 's threshold at ℓ_1 is also 2 since ℓ_1 can weakly accommodate f_5 alongside f_4 but not alongside $\{f_1, f_4\}$. Family f_6 's temporary threshold at ℓ_1 is 3 since ℓ_1 can weakly accommodate f_6 alongside $\{f_4, f_5\}$ but not alongside $\{f_1, f_4, f_5\}$. However, f_6 's threshold at ℓ_1 is 2 since f_6 's threshold at ℓ_1 cannot exceed f_5 's because f_5 has a higher priority and f_6 's temporary threshold is finite. Finally, f_7 's threshold at ℓ_1 is 0 since ℓ_1 cannot weakly accommodate f_7 alongside $\{f_4, f_5, f_6\}$. All three proposing families— f_4 , f_5 , and f_6 —have a threshold of 2 at ℓ_1 ; therefore, ℓ_1 tentatively accepts the two higher-priority proposing families— f_4 and f_5 —and permanently rejects the lowest-priority proposing family— f_6 . As ℓ_1 is able to accommodate $\{f_4, f_5, f_6\}$, one might be tempted to allow ℓ_1 to tentatively accept f_6 's proposal (as it would in the KDA algorithm). However, in this case, a choice function consistent with weak envy-freeness violates cardinal monotonicity. Suppose that f_1 , f_2 , f_4 , f_5 , and f_6 propose to ℓ_1 . Locality ℓ_1 cannot weakly accommodate f_4 alongside $\{f_1, f_2\}$, f_5 alongside $\{f_1, f_2, f_4\}$, and f_6 alongside $\{f_1, f_2, f_4, f_5\}$; therefore, weak envy-freeness dictates that ℓ_1 must only tentatively accept two families: f_1 and f_2 . However, cardinal monotonicity dictates that at most two families can be tentatively accepted when f_4 , f_5 , and f_6 propose.

Families f_1 and f_3 propose to ℓ_2 . Family f_5 's threshold at ℓ_2 is ∞ since f_5 has the highest priority and ℓ_2 can accommodate f_5 on its own. However, ℓ_2 cannot weakly accommodate f_1 alongside f_5 as this would require two units of d_2 and ℓ_2 has only one unit available; therefore, f_1 's threshold at ℓ_2 is 1. In contrast, f_2 's threshold at ℓ_2 is ∞ because ℓ_2 can weakly accommodate f_2 alongside $\{f_1, f_5\}$. (Recall from Algorithm 5 that if a family's threshold is ∞ , then it is allowed to exceed the thresholds of higher-priority families.) This situation illustrates how weak envy-freeness can improve efficiency over envy-freeness. Locality ℓ_2 cannot accommodate f_2 alongside $\{f_1, f_5\}$ as this would violate ℓ_2 's capacity for d_2 ; however, as f_2 does not require any unit of d_2 , ℓ_2 can weakly accommodate f_2 alongside $\{f_1, f_5\}$. The temporary threshold of f_7 at ℓ_2 is 3 since ℓ_2 can weakly accommodate f_7 alongside either one of $\{f_1, f_5\}$ or $\{f_1, f_2\}$, but not alongside $\{f_1, f_2, f_5\}$. However, since f_1 's threshold is 1, we set f_7 's threshold to 1 as well. The same reasoning applies to f_6 and f_3 while f_4 's threshold at ℓ_2 is 0 because ℓ_2 cannot weakly accommodate f_4 alongside $\{f_1, f_3\}$. Locality ℓ_2 tentatively accepts f_1 but permanently rejects f_3 .

No family proposes to ℓ_3 . The threshold of both f_4 and f_6 at ℓ_3 is ∞ since ℓ_3 can accommodate $\{f_4, f_6\}$. Family f_1 's threshold at ℓ_3 is 1 since ℓ_3 cannot weakly accommodate f_1 alongside f_4 . Finally, f_2 and f_7 propose to ℓ_4 . As f_7 has the highest priority for ℓ_4 and ℓ_4 can accommodate f_7 on its own, f_7 's threshold at ℓ_4 is ∞ , which means that ℓ_4 tentatively accepts f_7 . In contrast, ℓ_4 cannot weakly accommodate f_2 alongside f_7 ; therefore f_2 's threshold at ℓ_4 is 0 and ℓ_4 permanently rejects f_2 .

In Round 2, f_2 and f_3 both propose to ℓ_1 after having been permanently rejected by their respective first preferences (ℓ_4 and ℓ_2) in Round 1. As a result, f_4 's threshold at ℓ_1 rises to 3. This situation illustrates how a family's threshold can increase from one round to the next. In Round 1, f_4 's threshold at ℓ_1 is 2 because ℓ_1 cannot weakly accommodate f_4 alongside either one of $\{f_1, f_2\}$ or $\{f_1, f_3\}$. However, in Round 2, f_2 and f_3 propose to ℓ_1 but f_1 does not. As ℓ_1 can weakly accommodate f_4 alongside $\{f_2, f_3\}$, f_4 's threshold at ℓ_1 is 3. It follows that ℓ_1 continues to tentatively accept f_4 . In contrast, ℓ_4 cannot weakly

accommodate f_5 alongside $\{f_2, f_3, f_4\}$; therefore, f_5 's threshold at ℓ_1 is 0 and ℓ_1 permanently rejects f_5 . The third family that was permanently rejected in Round 1, f_6 , proposes to ℓ_2 in Round 2. As f_6 's threshold at ℓ_2 remains 1 and f_6 has the second-highest priority (after f_1) among proposing families, ℓ_2 permanently rejects f_6 .

In Round 3, f_5 proposes to ℓ_2 and is tentatively accepted since f_5 's threshold at ℓ_2 is ∞ . As a result, however, ℓ_2 permanently rejects f_1 which now has a threshold of 0. Family f_6 proposes to ℓ_3 and is also tentatively accepted since f_6 's threshold at ℓ_3 is ∞ . A consequence of f_6 's proposal to ℓ_3 is that f_1 's threshold at ℓ_3 rises to 2 because ℓ_3 can weakly accommodate f_1 alongside f_6 . Therefore, in Round 4, ℓ_3 tentatively accepts f_1 's proposal and the algorithm ends.

B.4 Threshold Knapsack Deferred Acceptance with Clinches

We use the example introduced in Online Appendix B.3 to illustrate the TKDAC algorithm. We show that

$$\mu^{TKDAC} = \begin{pmatrix} f_1 & f_2 & f_3 & f_4 & f_5 & f_6 & f_7 \\ \ell_3 & \ell_1 & \ell_1 & \ell_1 & \ell_2 & \ell_2 & \ell_4 \end{pmatrix}.$$

The only difference between μ^{TKDA} and μ^{TKDAC} is that $\mu^{TKDA}(f_6) = \ell_3$ while $\mu^{TKDAC}(f_6) = \ell_2$. Since $\ell_2 \succ_{f_6} \ell_3$, $\mu^{TKDAC} \succ \mu^{TKDA}$.

The three steps of the Clinching Round are displayed in Table 11. In Step 1(a), localities ℓ_1 , ℓ_2 and ℓ_3 do not reject any family since they can accommodate every family on its own. Locality ℓ_4 rejects every family that requires either two units of d_1 or one unit of d_2 , i.e., ℓ_4 rejects f_1 , f_4 , and f_5 . In Step 1(b), ℓ_1 is able to accommodate its three highest-priority families, which all receive a proposal. Locality ℓ_1 cannot weakly accommodate any other family alongside $\{f_1, f_2, f_3\}$ since every family takes up at least one unit of capacity in the first dimension; therefore none of the other families receive a proposal from ℓ_1 . Locality ℓ_2 can accommodate f_5 on its own but cannot weakly accommodate f_1 alongside f_5 since both families require a unit of d_2 and only one unit is available. In contrast, ℓ_2 can weakly accommodate f_2 alongside $\{f_1, f_5\}$ because f_2 does not require any units of d_2 . Therefore, ℓ_2 proposes to both f_5 and f_2 . None of the other families receive a proposal as this would violate the capacity of d_1 . Locality ℓ_3 is able to accommodate its two highest-priority families while ℓ_4 is only able to accommodate its highest-priority family. No other family receives a proposal from either ℓ_3 or ℓ_4 as this would violate their respective capacities in both dimensions. Family f_7 is the only family to receive a proposal from its first-preference locality, ℓ_4 . As a result, f_7 clinches ℓ_4 in Step 1(c). In Step 1(e), $\tilde{\nu}^2$ is constructed by giving f_7 the lowest priority for ℓ_1, ℓ_2, ℓ_3 . In particular, ℓ_2 's priority list is updated to

$$\tilde{\nu}_{\ell_2}^2 : f_5, f_1, f_2, f_6, f_3, f_4, f_7.$$

In Step 2(a), ℓ_4 rejects all families except f_7 because f_7 has clinched ℓ_4 and ℓ_4 cannot weakly accommodate any family alongside f_7 . In particular, ℓ_4 rejects its second-priority family f_2 . In Step 2(b), the same proposals occur as in Step 1; however, the fact that ℓ_4 has rejected f_2 means that ℓ_1 is now f_2 's most preferred locality that has not rejected f_2 . As ℓ_1 , proposes to f_2 , f_2 clinches ℓ_1 in Step 2(c). In Step 2(e), $\tilde{\nu}^3$ is constructed by moving f_2 to

Clinching Round - Step 1														
ℓ_1	(4, 2)			ℓ_2	(4, 1)			ℓ_3	(3, 1)			ℓ_4	(1, 0)	
f_1	(2, 1)	✓		f_5	(1, 1)	✓		f_4	(2, 0)	✓		f_7	(1, 0)	✓
f_2	(1, 0)	✓		f_1	(2, 1)			f_6	(1, 0)	✓		f_2	(1, 0)	
f_3	(1, 0)	✓		f_2	(1, 0)	✓		f_1	(2, 1)			\vdots		
f_4	(2, 0)			f_7	(1, 0)			\vdots						
f_5	(1, 1)			f_6	(1, 0)									
f_6	(1, 0)			f_3	(1, 0)									
f_7	(1, 0)			f_4	(2, 0)									

Clinching Round - Step 2														
ℓ_1	(4, 2)			ℓ_2	(4, 1)			ℓ_3	(3, 1)			ℓ_4	(1, 0)	
f_1	(2, 1)	✓		f_5	(1, 1)	✓		f_4	(2, 0)	✓		f_7	(1, 0)	✓
f_2	(1, 0)	✓		f_1	(2, 1)			f_6	(1, 0)	✓		f_2	(1, 0)	✗
f_3	(1, 0)	✓		f_2	(1, 0)	✓		f_1	(2, 1)			\vdots		
f_4	(2, 0)			f_6	(1, 0)			\vdots						
f_5	(1, 1)			f_3	(1, 0)									
f_6	(1, 0)			f_4	(2, 0)									
f_7	(1, 0)			f_7	(1, 0)									

Clinching Round - Step 3														
ℓ_1	(4, 2)			ℓ_2	(4, 1)			ℓ_3	(3, 1)			ℓ_4	(1, 0)	
f_1	(2, 1)	✓		f_5	(1, 1)	✓		f_4	(2, 0)	✓		f_7	(1, 0)	✓
f_2	(1, 0)	✓		f_1	(2, 1)			f_6	(1, 0)	✓		f_2	(1, 0)	✗
f_3	(1, 0)	✓		f_6	(1, 0)	✓		f_1	(2, 1)			\vdots		
f_4	(2, 0)			f_3	(1, 0)			\vdots						
f_5	(1, 1)			f_4	(2, 0)									
f_6	(1, 0)			f_2	(1, 0)									
f_7	(1, 0)			f_7	(1, 0)									

Table 11: Clinching Round of the TKDAC algorithm. Sizes and capacities in parentheses. ✓: a locality proposes to a family. ✗: a locality rejects a family. f_i : family f_i clinches a locality.

the bottom of ℓ_2 and ℓ_3 's priorities. f_2 's priority list is updated to

$$\tilde{\triangleright}_{\ell_2}^3 : f_5, f_1, f_6, f_3, f_4, f_2, f_7.$$

In Step 3(a), there are no new rejections as ℓ_1 can weakly accommodate any family alongside f_2 . In Step 3(b), there is one new proposal: ℓ_2 proposes to f_6 as ℓ_2 can weakly accommodate f_6 alongside $\{f_1, f_5\}$. This proposal was made possible by the fact that f_2 and f_7 have both moved below f_6 in ℓ_2 's priority list. However, ℓ_2 is f_6 's second preference and f_6 's first preference, ℓ_1 , has not rejected f_6 . It follows that no new clinch occurs in Step 3(c) so the Clinching Round ends in Step 3(d) and outputs $\tilde{\triangleright} = \tilde{\triangleright}^3$ such that

$$\begin{array}{ll} \triangleright_{\ell_1} : f_1, f_2, f_3, f_4, f_5, f_6, f_7 & \triangleright_{\ell_2} : f_5, f_1, f_6, f_3, f_4, f_2, f_7 \\ \triangleright_{\ell_3} : f_4, f_6, f_1, \dots & \triangleright_{\ell_4} : f_7, f_2, \dots \end{array}$$

With the priority profile $\tilde{\triangleright}$ constructed in the Clinching Round, the TKDA algorithm lasts four rounds. The first three rounds are displayed in Table 12. Changing the priority profile from \triangleright to $\tilde{\triangleright}$ affects the TKDA algorithm in one important way. With the constructed priority profile $\tilde{\triangleright}$, f_6 is the third highest-priority family at ℓ_2 . Since ℓ_2 can accommodate f_6 alongside $\{f_1, f_5\}$, f_6 's threshold at ℓ_2 is ∞ . It follows that ℓ_2 tentatively accepts ℓ_6 's proposal in Rounds 2-4, which is why $\mu^{TKDAC}(f_6) = \ell_2$.

When the TKDA algorithm is run with the true priority profile \triangleright , f_6 does not get a threshold of ∞ at ℓ_2 because f_2 and f_7 remain above f_6 on ℓ_2 's priority list. The Clinching Round's contribution in this example is to identify that f_2 and f_7 will necessarily be matched to a more preferred locality, i.e., f_2 will be matched to ℓ_1 (with $\ell_1 \succ_{f_2} \ell_2$) and f_7 will be matched to ℓ_4 (with $\ell_4 \succ_{f_7} \ell_2$). Thus, the Clinching Round identifies that f_2 and f_7 cannot cause a violation of weak envy-freeness or cardinal monotonicity when f_6 is matched to ℓ_2 .

TKDAC - Round 1

ℓ_1	(4, 2)	θ	ℓ_2	(4, 1)	θ	ℓ_3	(3, 1)	θ	ℓ_4	(1, 0)	θ
f_1	(2, 1)	∞	f_5	(1, 1)	∞	f_4	(2, 0)	∞	$\boxed{f_7}$	(1, 0)	∞
f_2	(1, 0)	∞	$\boxed{f_1}$	(2, 1)	1	f_6	(1, 0)	∞	$\boxed{\times_2}$	(1, 0)	0
f_3	(1, 0)	∞	f_6	(1, 0)	∞	f_1	(2, 1)	1	\vdots		
$\boxed{f_4}$	(2, 0)	2	$\boxed{\times_8}$	(1, 0)	1	\vdots					
$\boxed{f_5}$	(1, 1)	2	f_4	(2, 0)	0						
$\boxed{\times_6}$	(1, 0)	2	f_2	(1, 0)	0						
f_7	(1, 0)	0	f_7	(1, 0)	0						

TKDAC - Round 2

ℓ_1	(4, 2)	θ	ℓ_2	(4, 1)	θ	ℓ_3	(3, 1)	θ	ℓ_4	(1, 0)	θ
f_1	(2, 1)	∞	f_5	(1, 1)	∞	f_4	(2, 0)	∞	$\boxed{f_7}$	(1, 0)	∞
$\boxed{f_2}$	(1, 0)	∞	$\boxed{f_1}$	(2, 1)	1	f_6	(1, 0)	∞	f_2	(1, 0)	0
$\boxed{f_3}$	(1, 0)	∞	$\boxed{f_6}$	(1, 0)	∞	f_1	(2, 1)	1	\vdots		
$\boxed{f_4}$	(2, 0)	3	f_3	(1, 0)	1	\vdots					
$\boxed{\times_8}$	(1, 1)	0	f_4	(2, 0)	0						
f_6	(1, 0)	0	f_2	(1, 0)	0						
f_7	(1, 0)	0	f_7	(1, 0)	0						

TKDAC - Round 3

ℓ_1	(4, 2)	θ	ℓ_2	(4, 1)	θ	ℓ_3	(3, 1)	θ	ℓ_4	(1, 0)	θ
f_1	(2, 1)	∞	$\boxed{f_5}$	(1, 1)	∞	f_4	(2, 0)	∞	$\boxed{f_7}$	(1, 0)	∞
$\boxed{f_2}$	(1, 0)	∞	$\boxed{\times_1}$	(2, 1)	0	f_6	(1, 0)	∞	f_2	(1, 0)	0
$\boxed{f_3}$	(1, 0)	∞	$\boxed{f_6}$	(1, 0)	∞	f_1	(2, 1)	2	\vdots		
$\boxed{f_4}$	(2, 0)	3	f_3	(1, 0)	0	\vdots					
f_5	(1, 1)	0	f_4	(2, 0)	0						
f_6	(1, 0)	0	f_2	(1, 0)	0						
f_7	(1, 0)	0	f_7	(1, 0)	0						

Table 12: Rounds 1-3 of the TKDAC algorithm. Sizes and capacities in parentheses. $\boxed{f_i}$: f_i proposes and is tentatively accepted. $\boxed{\times_i}$: f_i proposes and is permanently rejected.

Preference type	Type 1	Type 2	Type 3	Type 4
Strong envy violations				
KTTC	590	1282	1302	4303
KDA	0	0	0	0
TKDA	0	0	0	0
Number of matched families				
KTTC	304	308	307	303
KDA	286	286	286	285
TKDA	236	230	230	232
Fraction of unfilled capacity				
KTTC	11.6%	12.0%	12.4%	13.0%
KDA	17.5%	15.5%	15.5%	17.3%
TKDA	26.5%	27.2%	27.3%	27.0%

Table 13: Outcomes of KTTC, KDA, and TKDA algorithms (three dimensions). Averages over 100 simulation rounds. Numbers in the top and middle panels are rounded to nearest integer; numbers in the bottom panel are rounded to 1 d.p.

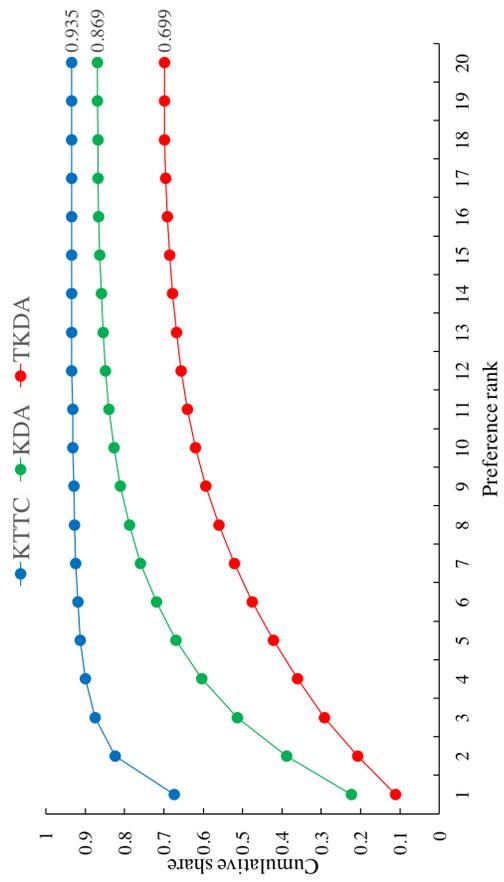
C Further simulation results

Table 13 and Figure 10 display the results for KTTC, KDA, and TKDA algorithms in the three-dimensional environment. The results are consistent with that of Table 7 across preference types and mechanisms. Figure 10 shows that the ranking of the mechanisms by efficiency in the one-dimensional setting (Figure 3) is preserved in the three-dimensional setting. Comparing Tables 7 and 13, one can see that mechanisms tend not to perform as well in the three-dimensional environment compared to the one-dimensional environment: the KTTC mechanism generates more strong envy violations and all three mechanisms are less efficient. These findings are intuitive in the sense that having three dimensions increases the heterogeneity among families' sizes. Therefore, the KTTC mechanism has more difficulty finding Pareto-improving trades and those trades are more likely to create strong envy violations; moreover, in the KDA and TKDA mechanism, localities may have to reject more families in order to prevent violations of strong envy.

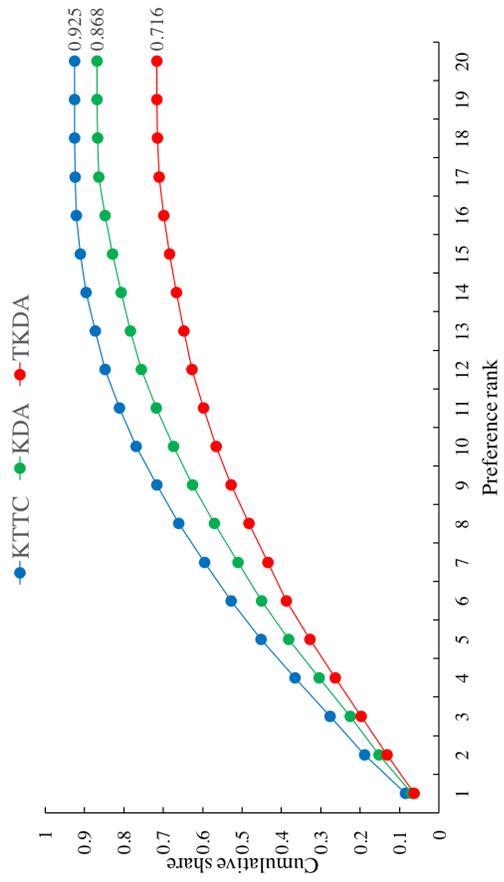
D Relationships to prior models

Our model generalizes a number of existing matching models, including the following:

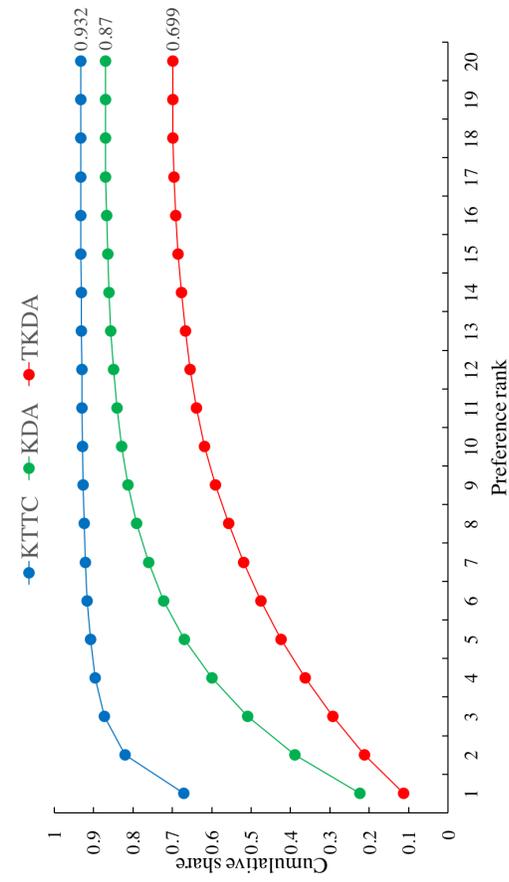
- School choice (Abdulkadiroğlu and Sönmez, 2003): Every student takes up a single seat at any school. Let us relabel a student as a family and a school as a locality. In our model, this corresponds to having only one dimension d ($|D| = 1$) and any family having a size of one ($\nu_d^f = 1$ for all $f \in F$).
- Controlled school choice or college admissions with affirmative action and m type-specific quotas (Abdulkadiroğlu and Sönmez, 2003; Abdulkadiroğlu, 2005; Westkamp, 2013): Each student is one of m types and each school has a quota for each of the m types. Let us again relabel a student as a family, a school as a locality and a type



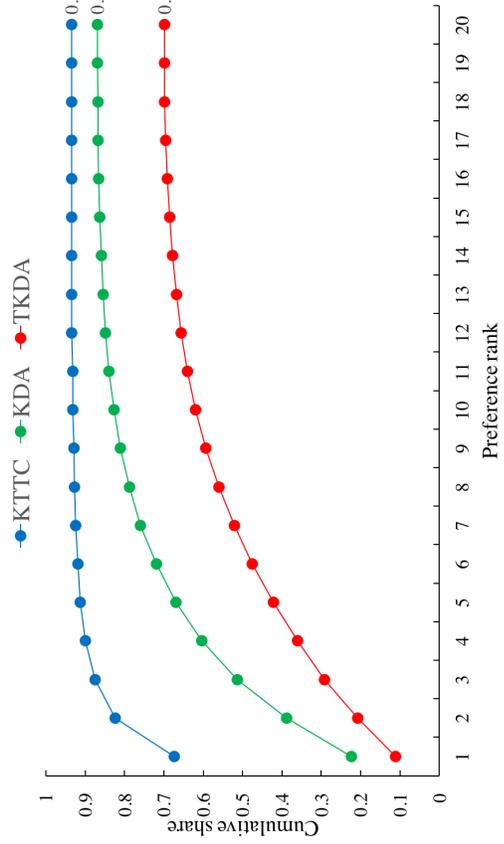
(a) Type 1



(b) Type 2



(c) Type 3



(d) Type 4

Figure 10: Preference distributions for outcomes under different matching mechanisms (three dimensions). Labelled numbers: fractions of matched families).

as a dimension. In our model, this corresponds to having m dimensions ($|D| = m$). Each family takes up one unit of capacity in exactly one of the dimensions (ν^f is a m -dimensional unit vectors for every $f \in F$).

- School choice with majority quotas (Kojima, 2012; Hafalir et al., 2013): Each student is either a majority or a minority student. Each school has a overall cap on the number of students, which includes a cap for majority students. Let us again relabel a majority/minority student as a majority/minority family and a school as a locality. Let us also relabel “any student seats” as dimension d_1 and “majority student seats” as dimension d_2 ($|D| = 2$). In our model, the capacity of any locality for d_1 is greater than the capacity for d_2 ($\kappa_{d_1}^\ell > \kappa_{d_2}^\ell$ for all $\ell \in L$). A majority family f takes up one unit of capacity in each dimension ($\nu^f = (1, 1)$) whereas a minority family f' only takes up a unit of capacity in d_1 ($\nu^{f'} = (1, 0)$).
- Hungarian college admissions (Biró et al., 2010): Students take up a college seat as well as a faculty seat. Both colleges and faculties have their own capacities. Let us relabel a student as a family and a college as a locality. Let us also relabel “college capacity” as the capacity of the locality in dimension d_1 ($\kappa_{d_1}^\ell$). Let us relabel the faculties as the remaining dimensions $D \setminus \{d_1\}$. Therefore, each family f 's size is $\nu^f = (1, 0, 0, \dots, 1, \dots, 0, 0)$ where the second “1” is the unit of capacity taken up in $d \in D \setminus \{d_1\}$.
- Allocation of trainee teachers to schools in Slovakia and Czechia (Cechlárová et al., 2015): Teachers are required to teach two out of three subjects and each school has a capacity for all three subjects. Let us relabel a teacher as a family, a school as a locality, and a subject as a dimension. In our model, this corresponds to having three dimensions ($|D| = 3$) and the size of every family f being either $\nu^f = (0, 1, 1)$, $\nu^f = (1, 0, 1)$, or $\nu^f = (1, 1, 0)$.
- College admission with multidimensional privileges in Brazil (Aygün and Bó, 2020): Students can claim any combination of three privileges. Colleges have quotas for each privilege, but a single student can claim more than one privilege. Let us relabel a student as a family, a college as a locality, and a privilege as a dimension. In our model, this corresponds to having three dimensions ($|D| = 3$) and the size of every family being an element of $\{0, 1\}^3$.
- Object allocation (Nguyen et al., 2016) or course assignment (Budish, 2011): Agents (students) demand a certain number of different objects (courses) that are supplied by a single seller (a business school). Let us relabel agents (students) as families, different objects (courses) as dimensions, and the single seller (the business school) as a single locality. In our model, this corresponds to having only one locality ($|L| = 1$).
- Resident-hospital matching with sizes (McDermid and Manlove, 2010): Doctors apply to hospitals, but the doctors can take up more than one seat at a hospital, e.g., because they arrive as couples. Let us relabel doctors as families and hospitals as localities. In

our model, this corresponds to having one dimension ($|D| = 1$) and families having an arbitrary size.⁴⁴

Most of the models described above use further assumptions and develop solution approaches that suit their particular contexts but differ substantially from ours. Nevertheless, as we note throughout the paper, several impossibility and complexity results established in these papers will apply immediately to our framework.

⁴⁴This model in turn generalizes the resident-hospital matching with inseparable couples (i.e., when couples have the same preference list and prefer to be unmatched to being in different hospitals) as well as resident-hospital matching with couples which have “consistent” preferences (McDermid and Manlove, 2010, Lemma 2.1). In both cases, we set $|D| = 1$ and $\nu_d^f \in \{1, 2\}$ in our model.