MATCHING MECHANISMS FOR REFUGEE RESETTLEMENT

DAVID DELACRÉTAZ, SCOTT DUKE KOMINERS, AND ALEXANDER TEYTELBOYM

ONLINE APPENDIX

Contents

Appendix A. Proofs of the Results Stated in the Main Text

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 $\mu^{\text{KTTC}}(\ell)$. This holds trivially for the null. For any non-null locality $\ell \neq \emptyset$, in any Round $i \geq 1$, 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 $\mu^1(\ell) = \varnothing$, 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 $\mu^{\text{KTTC}}(\ell)$. We next show that the KTTC mechanism is Pareto-efficient (PE) and strategy-proof (SP).

Proof of (PE). For each Round $i = 1, \ldots, N$, let F^i be the set of families that have been permanently matched by the end of Round $i-1$. (Note that $F^1 = \varnothing$ and $F^{N+1} = F$.) If the algorithm ends in Round N by matching all remaining families to the null, we use the convention that $\mu^{N+1} = \mu^N$.

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 $F^1 = \varnothing$. We now show that if our induction hypothesis holds for some $i = 1, \ldots, N$, 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 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 of the fact that for all $f \in F^i$, $\mu^i(f) = \mu^{i+1}(f)$.

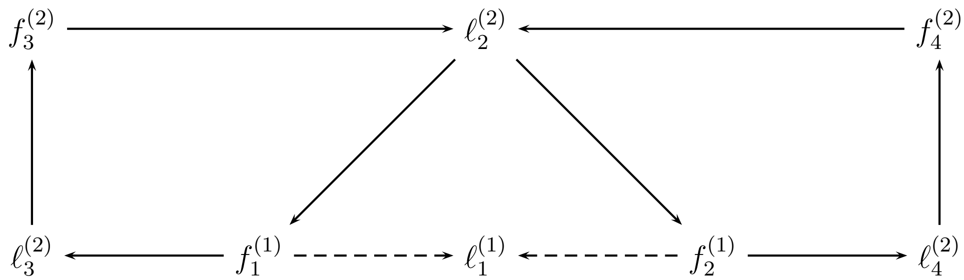

FIGURE A.1. Counterexample for the proof of Theorem 1. $f \to \ell$: ℓ is f's first choice. $f \dashrightarrow \ell$: ℓ is f's second choice and f prefers ℓ to its endowment. $\ell \to f$: ℓ is f's endowment. Superscripts denote sizes and capacities, respectively.

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 $\mu^{N+1} = \mu^{\text{KTTC}}$ 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 $\mu^{\text{KTTCE}} = \mu^{\text{KTTC}}$ when every family is endowed with $\emptyset$ (Proposition 2). □

Proof of Theorem 1. The proof is by counterexample. There are four families, four localities, and one dimension. The endowment is

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

The preferences of families are as follows:

$$\succ_{f_1}: \ell_3, \ell_1, \boldsymbol{\ell_2}, \dots \qquad \succ_{f_2}: \ell_4, \ell_1, \boldsymbol{\ell_2}, \dots \qquad \succ_{f_3}: \ell_2, \boldsymbol{\ell_3}, \dots \qquad \succ_{f_4}: \ell_2, \boldsymbol{\ell_4}, \dots,$$

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

$$\nu = \begin{array}{c} \\ d_1 \end{array} \begin{array}{c} \begin{array}{cccc} f_1 & f_2 & f_3 & f_4 \end{array} \\ \begin{pmatrix} 1 & 1 & 2 & 2 \end{pmatrix} \end{array} \qquad \kappa = \begin{array}{c} \\ d_1 \end{array} \begin{array}{c} \begin{array}{cccc} \ell_1 & \ell_2 & \ell_3 & \ell_4 \end{array} \\ \begin{pmatrix} 1 & 2 & 2 & 2 \end{pmatrix} \end{array}.$$

For ease of exposition, we illustrate this counterexample in Figure A.1.

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

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

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 $f_1 \to \ell_3 \to f_3 \to \ell_2$. 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., reporting $\succ'_{f_1}: \ell_3, \boldsymbol{\ell_2}, \ldots$. Using analogous reasoning as for the true preference profile, it is easy to check that μ' is the 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., $\succ'_{f_2}: \ell_4, \boldsymbol{\ell_2}, \ldots$.

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 Theorem 2. The proof is by counterexample. There are four families, three localities, and two dimensions. The endowment is

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

The preferences of families are as follows:

$$\succ_{f_1}: \ell_3, \boldsymbol{\ell_2}, \ldots \quad \succ_{f_2}: \ell_4, \boldsymbol{\ell_2}, \ldots \quad \succ_{f_3}: \ell_2, \ell_4, \boldsymbol{\ell_3}, \ldots \quad \succ_{f_4}: \ell_2, \ell_3, \boldsymbol{\ell_4}, \ldots,$$

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

$$\nu = \begin{array}{c} \\ d_1 \\ d_2 \end{array} \begin{array}{c} \begin{array}{cccc} f_1 & f_2 & f_3 & f_4 \end{array} \\ \begin{pmatrix} 2 & 1 & 2 & 3 \\ 1 & 2 & 3 & 2 \end{pmatrix} \end{array} \qquad \kappa = \begin{array}{c} \\ d_1 \\ d_2 \end{array} \begin{array}{c} \begin{array}{ccc} \ell_2 & \ell_3 & \ell_4 \end{array} \\ \begin{pmatrix} 4 & 3 & 3 \\ 4 & 3 & 3 \end{pmatrix} \end{array}.$$

For ease of exposition, we illustrate this counterexample in Figure A.2.

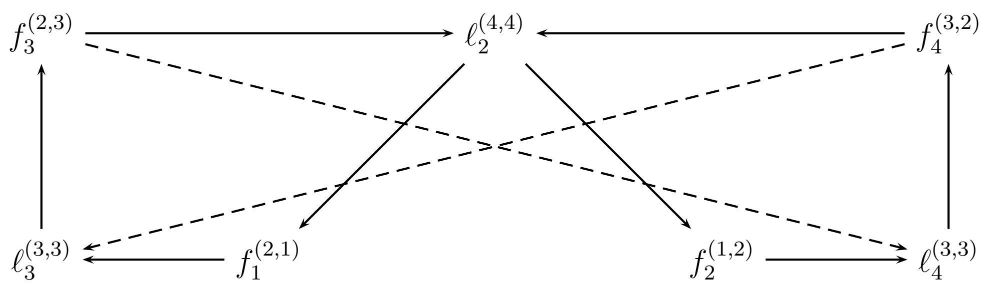

FIGURE A.2. Counterexample for the proof of Theorem 2. $f \to \ell$: ℓ is f's first choice. $f \dashrightarrow \ell$: ℓ is f's second choice and f prefers ℓ to its endowment. $\ell \to f$: ℓ is f's endowment. Superscripts denote sizes and capacities, respectively.

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

$$(f_2, \ell_4, f_4, \ell_3, f_3, \ell_2), \quad (f_1, \ell_3, f_3, \ell_4, f_4, \ell_2), \quad \text{and} \quad (f_3, \ell_4, f_4, \ell_3).$$

Executing these chains yields the following three matchings, respectively:

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

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 $\tilde{\mu}(f_3) = \ell_3$. Then, neither f_1 nor f_4 can be matched to ℓ_3; hence $\tilde{\mu}(f_1) = \ell_2$ and $\tilde{\mu}(f_4) = \ell_4$. Therefore, we must have that $\tilde{\mu}(f_2) = \ell_2$. It follows that $\tilde{\mu} = \mu^E$, a contradiction.

Second, consider the case where $\tilde{\mu}(f_3) = \ell_2$. Then, f_2 cannot be matched to ℓ_2 so $\tilde{\mu}(f_2) = \ell_4$. As a result, f_4 cannot be matched to either ℓ_2 or ℓ_4; hence, $\tilde{\mu}(f_4) = \ell_3$. It follows that f_1 cannot be matched to ℓ_3; therefore $\tilde{\mu}(f_3) = \ell_2$ and $\tilde{\mu} = \mu$, a contradiction.

Third, consider the case where $\tilde{\mu}(f_3) = \ell_4$. Neither f_2 nor f_4 can be matched to ℓ_4; therefore, $\tilde{\mu}(f_2) = \ell_2$ and we must have that either $\tilde{\mu}(f_4) = \ell_2$ or $\tilde{\mu}(f_4) = \ell_3$. If $\tilde{\mu}(f_4) = \ell_2$, 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 $\succ'_{f_3}$: $\ell_2, \boldsymbol{\ell_3}, \ldots$; 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 $\succ'_{f_4}$: $\ell_2, \boldsymbol{\ell_4}, \ldots$, 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 $\mu^{\text{KTTCE}}(\ell)$. This holds trivially for the null. For a non-null locality $\ell \neq \emptyset$, in any Round i, if there exists a family $f \in \mu^{i+1}(\ell) \setminus \mu^i(\ell)$, then f moves to ℓ as part of a feasible cycle, and therefore ℓ can accommodate f alongside the families that remain at ℓ after the cycle is carried out. If follows that 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, ℓ can accommodate $\mu^{\text{KTTCE}}(\ell)$.

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 $\mu^{\text{KTTCE}}(f) \succeq_f \ell$. 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 $\mu^{\text{KTTCE}}(f)$. Therefore, $\mu^{\text{KTTCE}}(f) \succeq_f \ell$, 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 and m' be the rounds in which f gets permanently matched with reports $\succ_f$ and $\succ'_f$, respectively.

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, i.e., $\ell, \ell' \in L_f^m$. 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: $m > m'$. If f reports $\succ'_f$, f points at and is permanently matched to ℓ' in Round m'. If $\ell' = \emptyset$, then $\ell \succeq_f \ell'$ since every family prefers every locality to the null. The remainder of the proof assumes that $\ell' \neq \emptyset$. In that case, if f reports $\succ'_f$, a feasible cycle $f \to \ell' \to f_2 \to \ell_2 \to \ldots \to f_n \to \ell_n \to f$ appears in Round m' (with $n = 1, 2, \ldots$).

Suppose now that f reports $\succ_f$. For any $i = m', \ldots, m$, denote by L_f^i the set of localities that have not permanently rejected f at the start of Round i. We proceed by induction with the following hypothesis: In Round $i = m', \ldots, m$, $\ell' \in L_f^i$, ℓ' points at f_2, each f_j with $j = 2, \ldots, n$ points at ℓ_j, each ℓ_j with $j = 2, \ldots, n-1$ points at f_{j+1}, and ℓ_n points at f.

We first show that our inductive hypothesis holds for $i = m'$. By construction, how family f points does not affect how other families point and which families each locality rejects until after f is permanently matched. Therefore, irrespective of whether f reports $\succ_f$ or $\succ'_f$, ℓ' has not rejected f at the start of Round m' (hence $\ell' \in L_f^{m'}$) and, in Round m', ℓ' points at f_2, each f_j with $j = 2, \ldots, n$ points at ℓ_j, each ℓ_j with $j = 2, \ldots, n-1$ points at f_{j+1}, and ℓ_n points at f.

We now suppose that our induction hypothesis holds for some $i = m', \ldots, m-1$ and show that it holds for $i+1$. As f is not matched until Round $m > i$ and each family and locality can be part of at most one cycle, none of $f_2, \ldots, f_n$ and $\ell', \ell_2, \ldots, \ell_n$ are in a cycle in Round m', hence none of them are permanently matched in Round i. Moreover, as the cycle $f \to \ell' \to f_2 \to \ell_2 \to \ldots \to f_n \to \ell_n \to f$ is feasible, each ℓ_j with $j = 2, \ldots, n-1$ can accommodate f_j alongside $\mu^i(\ell_j) \setminus \{f_{j+1}\}$. Therefore, ℓ_j does not permanently reject f_j in Round i. Analogously, ℓ' does not permanently reject f in Round i (hence $\ell' \in L_f^{i+1}$) and ℓ_n does not permanently reject f_n in Round i. It follows that, in Round $i+1$, ℓ' continues to point at f_2, each f_j with $j = 2, \ldots, n$ continues to point at ℓ_j, each ℓ_j with $j = 2, \ldots, n-1$ continues to point at f_{j+1}, and ℓ_n continues to point at f.

By induction, we conclude that $\ell' \in L_f^m$. Hence, since ℓ is f's most preferred family in L_f^m by construction, we have that $\ell \succeq_f \ell'$. □

Proof of Proposition 2. We first show that all cycles that appear in the KTTCE algorithm are feasible. Towards a contradiction, suppose that an infeasible cycle $f_1 \to \ell_1 \to f_2 \to \ell_2 \to \ldots \to f_n \to \ell_n \to f_1$ appears in some Round i. Then, there exists $j = 1, \ldots, n$ such that ℓ_j cannot accommodate f_j alongside $\mu^i(\ell_j) \setminus \{f_{j+1}\}$ (letting $f_{n+1} = f_1$). As the null does not point (and therefore is not in any cycles), where $\ell_j \neq \emptyset$, so $\mu^E(\ell_j) = \varnothing$; 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 $\mu^E(\ell_j) = \varnothing$, 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 shown that all cycles that appear in the KTTCE algorithm are feasible, which implies that the KTTCE algorithm does not enter the Rejection Stage in any round. As a result, each round of the KTTCE algorithm coincides with the corresponding round of the KTTC algorithm. □

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., $\mu^{\text{KTTCE}} = \mu^E$. 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: $\mu^E = \mu^1 = \mu^2 = \ldots = \mu^N = \mu^{N+1} = \mu^{\text{KTTCE}}$. (If the algorithm ends in Round N by matching all remaining families to the null, we use the convention that $\mu^{N+1} = \mu^N$.) 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 $M^* = \varnothing$; therefore it remains to show that $M^* = \varnothing$.

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 $i = 1, \ldots, N$, 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 that f 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. (As the null never permanently rejects any family, $\ell \neq \emptyset$.) 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). Since a family is never permanently rejected by its endowment, we have that $\ell \neq \mu^E(f)$. By construction, f' has the highest priority at ℓ among all families that are not permanently matched. 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 $\mu^E(\ell)$ have been 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 $\mu^E(\ell)$. As all families in $\mu^E(\ell)$ are permanently matched to ℓ at the start of Round i, 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 $\tilde{f} \rhd_\ell f'$, 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 $f' \rhd_\ell \tilde{f}$ or $f' = f$. 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 $M^* = \varnothing$, as desired. $\square$

Proof of Theorem 5. Denote by N the last round of the KDA algorithm. Now, for any family $f \in F$, any locality $\ell \in L$, and any Round $i = 1, \ldots, N$, let $\widehat{R}_i^f(\ell)$ be the set of families that have a higher priority than f at ℓ and propose in Round i to either ℓ or a less preferred locality (because they have already been permanently rejected by ℓ in a previous round).

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

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

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

We next show that μ^{KDA} is interference-free. Consider any family f; we need to show that f does not interfere with μ. By construction, f proposes to and is tentatively accepted by $\mu^{\text{KDA}}(f)$ in the last round; therefore, $\mu(f)$ can weakly accommodate f alongside $\widehat{R}_N^f(\mu^{\text{KDA}}(f))$. Again by construction, the families in $\widehat{R}_N^f(\mu^{\text{KDA}}(f))$ are exactly those that end up matched to either $\mu^{\text{KDA}}(f)$ or a less preferred locality. It follows that $\widehat{R}_N^f(\mu^{\text{KDA}}(f)) = \widehat{F}_{\mu^{\text{KDA}}}^f$, which implies that f does not interfere with μ^{KDA}.

We finally show that μ^{KDA} dominates all other interference-free matchings. Towards a contradiction, suppose that there exists an interference-free matching μ such that $\mu(f_1) \succ_{f_1} \mu^{\text{KDA}}(f_1)$ 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 $\widehat{R}_{i_n}^{f_n}(\mu(f_n))$. If all families in $\widehat{R}_{i_n}^{f_n}(\mu(f_n))$ were matched at μ to either $\mu(f_n)$ or a less preferred locality, then f_n would interfere with μ, a contradiction. Therefore, there exists a family $f_{n+1} \in \widehat{R}_{i_n}^{f_n}(\mu(f_n))$ such that $\mu(f_{n+1}) \succ_{f_{n+1}} \mu(f_n)$. By construction, as f_{n+1} proposes in Round i_n of the KDA algorithm to either $\mu(f_n)$ or a less preferred locality, $\mu(f_{n+1})$ 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 $\mu^{\mathrm{TKDA}}(\ell)$. Consider a locality $\ell \in L$. By construction, all families in $\mu^{\mathrm{TKDA}}(\ell)$ propose to and are tentatively accepted by ℓ in the last round of the TKDA algorithm. Hence, for any $f \in \mu^{\mathrm{TKDA}}(\ell)$,

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

If ℓ cannot weakly accommodate f alongside $\mu^{\mathrm{TKDA}}(\ell) \cap \widehat{F}_\ell^f$, 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 $f \in \mu^{\mathrm{TKDA}}(\ell)$ alongside $\mu^{\mathrm{TKDA}}(\ell) \cap \widehat{F}_\ell^f$. As we showed in the proof of Theorem 5, this implies that ℓ can accommodate $\mu^{\mathrm{TKDA}}(\ell)$.

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

Proof that 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 each $f \in F$, $\widehat{G}_\ell^f = G \cap \widehat{F}_\ell^f$ to be the families in G that have a higher priority at 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, $C_\ell : 2^F \to 2^F$, 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, $C_\ell(G)$ is a function that selects 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 yet permanently rejected them. Localities tentatively accept or permanently reject proposals based on their choice functions, i.e., if locality ℓ receives proposals from families in G, then families in $C_\ell(G)$ are tentatively accepted and families in $G \backslash C_\ell(G)$ are permanently rejected.

Theorems 3 and 11 of Hatfield and Milgrom (2005) imply that the DA mechanism is strategy-proof for families if the choice function of every locality satisfies 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 mechanism is strategy-proof, it remains to prove that the choice function of every locality satisfies the (S) and (CM) conditions.

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

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

(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 $g \rhd_\ell f$ 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 A.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 A.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 A.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 $H = G \cup \{f_1, \ldots, f_m\}$.

Claim A.1. For each $i = 1, \ldots, m$, $|C_\ell(G \cup \{f_1, \ldots, f_{i-1}\})| \leq |C_\ell(G \cup \{f_1, \ldots, f_{i-1}, f_i\})|$.

Claim A.1 implies that

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

which implies the desired result. Therefore, it remains to prove Claim A.1.

Proof of Claim A.1. Fix some $i \leq m$ and define $G' = G \cup \{f_1, \ldots, f_{i-1}\}$ and $H' = \{f_1, \ldots, f_{i-1}, f_i\}$. 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

$$|C_\ell(H') \setminus C_\ell(G')| \geq |C_\ell(G') \setminus C_\ell(H')|. \tag{1}$$

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, $C_\ell(H') \setminus C_\ell(G') = \varnothing$; hence, by inequality (1), we need to show that $C_\ell(G') \setminus C_\ell(H') = \varnothing$. 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 C_ℓ, we have that

$$\theta_\ell^g(G') \geq |\widehat{G}_\ell'^g| + 1; \tag{2}$$

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

$$\theta_\ell^g(H') < |\widehat{H}_\ell'^g| + 1. \tag{3}$$

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

Sub-case 1.1: $g \rhd_\ell f_i$. In this case, as $H' = G' \cup \{f_i\}$, $\widehat{G}_\ell'^g = \widehat{H}_\ell'^g$ so $|\widehat{G}_\ell'^g| = |\widehat{H}_\ell'^g|$ 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 (2) and (3) yields

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

a contradiction.

Sub-case 1.2: $f_i \rhd_\ell g$. On the one hand, inequality (3) implies that $\theta_\ell^g(H') \neq \infty$, so we can apply Lemma A.1(i) to obtain that $\theta_\ell^g(G') \neq \infty$ and Lemma A.1(iii) to obtain that

$$\theta_\ell^{f_i}(G') \geq \theta_\ell^g(G'). \tag{4}$$

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

$$\theta_\ell^{f_i}(H') < |\widehat{H}_\ell'^{f_i}| + 1. \tag{5}$$

As $H' = G' \cup \{f_i\}$, $\widehat{G}_\ell'^{f_i} = \widehat{H}_\ell'^{f_i}$ so $|\widehat{G}_\ell'^{f_i}| = |\widehat{H}_\ell'^{f_i}|$ 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 (5) yields

$$\theta_\ell^{f_i}(G') = \theta_\ell^{f_i}(H') < |\widehat{H}_\ell'^{f_i}| + 1 = |\widehat{G}_\ell'^{f_i}| + 1. \tag{6}$$

Combining inequality (6) with the fact that $|\widehat{G}'^{f_i}_\ell| \leq |\widehat{G}'^{g}_\ell|$ (as $f_i \rhd g$) and inequality (2) yields

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

We conclude that $\theta^{f_i}_\ell(G') < \theta^{g}_\ell(G')$, which contradicts inequality (4).

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 (1), 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 $g_1 \rhd_\ell g_2$. Since $g_1, g_2 \in C_\ell(G')$, by the definition of the choice function, we have that

$$\theta^{g_1}_\ell(G') \geq |\widehat{G}'^{g_1}_\ell| + 1 \quad \text{and} \quad \theta^{g_2}_\ell(G') \geq |\widehat{G}'^{g_2}_\ell| + 1, \tag{7}$$

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

$$\theta^{g_1}_\ell(H') < |\widehat{H}'^{g_1}_\ell| + 1 \quad \text{and} \quad \theta^{g_2}_\ell(H') < |\widehat{H}'^{g_2}_\ell| + 1. \tag{8}$$

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

Sub-case 2.1: $g_1 \rhd_\ell f_i$. In this case, as $H' = G' \cup \{f_i\}$, $\widehat{G}'^{g_1}_\ell = \widehat{H}'^{g_1}_\ell$ so $|\widehat{G}'^{g_1}_\ell| = |\widehat{H}'^{g_1}_\ell|$ and, as a family's threshold only depends on higher-priority families (Algorithm 5), $\theta^{g_1}_\ell(G') = \theta^{g_1}_\ell(H')$. Combining these observations with inequalities (7) and (8) yields

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

a contradiction.

Sub-case 2.2: $f_i \rhd_\ell g_1$. Inequality (8) implies that $\theta^{g_2}_\ell(H') \neq \infty$ so we can apply Lemma A.1(i) to obtain that $\theta^{g_2}_\ell(G') \neq \infty$ and Lemma A.1(iii) to obtain that $\theta^{g_1}_\ell(G') \geq \theta^{g_2}_\ell(G')$. Moreover, as $g_1 \rhd_\ell g_2$ and $g_1 \in G$, we have that $\widehat{G}^{g_1}_\ell \subset \widehat{G}^{g_2}_\ell$ so $|\widehat{G}^{g_1}_\ell| < |\widehat{G}^{g_2}_\ell|$. Combining these observations with inequality (7) yields

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

We therefore conclude that

$$\theta^{g_1}_\ell(G') \geq |\widehat{G}'^{g_1}_\ell| + 2. \tag{9}$$

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

Sub-sub-case 2.2.1: $\theta^{g_1}_\ell(H') = 0$. 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 $h \rhd_\ell g_1$ and (ii) ℓ cannot weakly accommodate h alongside $\widehat{H}'^{h}_\ell$. First, as h has a weakly higher priority than g_1 and $\theta^{g_1}_\ell(G') \neq \infty$ (by inequality (8) and Lemma A.1(i)), we can apply Lemma A.1(iii) to obtain that $\theta^{g_1}_\ell(G') \leq \theta^{h}_\ell(G')$. Second, as $G' \subseteq H'$, we have that $\widehat{G}'^{h}_\ell \subseteq \widehat{H}'^{h}_\ell$; therefore, $\widehat{H}'^{h}_\ell$ is a subset of $\widehat{F}^{h}_\ell$ that contains all families in $\widehat{G}'^{h}_\ell$ and alongside which ℓ cannot weakly accommodate h. By the definition of thresholds (Algorithm 5), it follows that $\theta^{h}_\ell(G') \leq |\widehat{H}'^{h}_\ell|$. Third, as

$H' = G' \cup \{f_i\}$, we have that $|\widehat{H}'^h_\ell| \leq |\widehat{G}'^h_\ell| + 1$. Fourth, as h has a weakly higher priority than g_1, we have that $|\widehat{G}'^h_\ell| \leq |\widehat{G}'^{g_1}_\ell|$. Combining these four observations yields

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

which contradicts inequality (9).

Sub-sub-case 2.2.2: $\theta^{g_1}_\ell(H') \in \mathbb{Z}_{>0}$. In this case, we can apply Lemma A.1(ii) to obtain that $\theta^{g_1}_\ell(H') \geq \theta^{g_1}_\ell(G')$. Moreover, as $H' = G' \cup \{f_i\}$ and $f_i \in \widehat{H}'^{g_1}_\ell$ (by the assumption of Sub-case 2.2), $f_i \rhd_\ell g_1$), we have that $|\widehat{H}'^{g_1}_\ell| = |\widehat{G}'^{g_1}_\ell| + 1$. Combining these observations with inequality (9) yields

$$\theta^{g_1}_\ell(H') \geq \theta^{g_1}_\ell(G') \geq |G'^{g_1}_\ell| + 2 = |H'^{g_1}_\ell| + 1,$$

which contradicts inequality (8) and completes the proof of Claim A.1. □

Proof that TKDA is interference-free.

Let N be the number of rounds after which the TKDA algorithm terminates. For each family $f \in F$, each locality $\ell \in L$, and each Round $i = 1, \ldots, N$, denote by Π^i_ℓ the set of families that propose to ℓ and by $\theta^f_\ell(\Pi^i_\ell)$ the threshold of f for ℓ in Round i.

Fix a family $f \in F$ and, for notational simplicity, let $\ell = \mu^{\text{TKDA}}(f)$. We need to show that f does not interfere with μ^{TKDA}, i.e., that ℓ can weakly accommodate f alongside $\widehat{F}^f_{\mu^{\text{TKDA}}}$.

If $\theta^f_\ell(\Pi^N_\ell) = \infty$, then by definition ℓ can weakly accommodate f alongside $\widehat{F}^f_\ell$, hence alongside $\widehat{F}^f_{\mu^{\text{TKDA}}}$. The remainder of the proof considers the case in which $\theta^f_\ell(\Pi^N_\ell) \neq \infty$.

By construction, ℓ tentatively accepts f in Round N. By definition, it follows that $\theta^f_\ell(\Pi^N_\ell) > 0$; hence, ℓ can weakly accommodate f alongside $\Pi^N_\ell \cap \widehat{F}^f_\ell$. We complete the proof by showing that $(\Pi^N_\ell \cap \widehat{F}^f_\ell) = \widehat{F}^f_{\mu^{\text{TKDA}}}$. Observing that, by construction, $(\Pi^N_\ell \cap \widehat{F}^f_\ell) \subseteq \widehat{F}^f_{\mu^{\text{TKDA}}}$, it remains to show that the case where $(\Pi^N_\ell \cap \widehat{F}^f_\ell) \subset \widehat{F}^f_{\mu^{\text{TKDA}}}$ leads to a contradiction.

Towards a contradiction, suppose that $(\Pi^N_\ell \cap \widehat{F}^f_\ell) \subset \widehat{F}^f_{\mu^{\text{TKDA}}}$. Let g be the highest-priority family in $\widehat{F}^f_{\mu^{\text{TKDA}}} \setminus (\Pi^N_\ell \cap \widehat{F}^f_\ell)$; g is then the highest-priority family to be permanently rejected by ℓ in the TKDA algorithm. Let i be the round in which that permanent rejection occurs. Then, g's threshold in Round i is smaller than its priority rank. As a family's threshold only depends on higher-priority families, we have that

$$\theta^g_\ell(\Pi^i_\ell \cap \widehat{F}^g_\ell) = \theta^g_\ell(\Pi^i_\ell) < |\Pi^i_\ell \cap \widehat{F}^g_\ell| + 1. \tag{10}$$

As g is the highest-priority family to be permanently rejected by ℓ, every family with a higher priority than g that proposes to ℓ in Round i also proposes to ℓ in Round N, hence $(\Pi^i_\ell \cap \widehat{F}^g_\ell) \subseteq (\Pi^N_\ell \cap \widehat{F}^g_\ell)$. Moreover, (10) implies that $\theta^g_\ell(\Pi^i_\ell \cap \widehat{F}^g_\ell) \neq \infty$, which by Lemma A.1(i) yields $\theta^g_\ell(\Pi^N_\ell \cap \widehat{F}^g_\ell) \neq \infty$. Therefore, Lemma A.1(ii) applies and we have that

$$\theta^g_\ell(\Pi^N_\ell \cap \widehat{F}^g_\ell) \leq \theta^g_\ell(\Pi^i_\ell \cap \widehat{F}^g_\ell) + |\Pi^N_\ell \cap \widehat{F}^g_\ell| - |\Pi^i_\ell \cap \widehat{F}^g_\ell|. \tag{11}$$

Combining (10) and (11) yields

$$\theta^g_\ell(\Pi^N_\ell \cap \widehat{F}^g_\ell) < |\Pi^N_\ell \cap \widehat{F}^g_\ell| + 1. \tag{12}$$

By assumption (and as a family's threshold only depends on higher-priority families), we have that $\theta^f_\ell(\Pi^N_\ell \cap \widehat{F}^f_\ell) = \theta^f_\ell(\Pi^N_\ell) \neq \infty$. Then, as $g \rhd_\ell f$, Lemma A.1(iii) applies and yields $\theta^f_\ell(\Pi^N_\ell \cap \widehat{F}^f_\ell) \leq \theta^g_\ell(\Pi^N_\ell \cap \widehat{F}^g_\ell)$. Combining that inequality with (12), using again the property that a family's threshold only depends on higher-priority families, and observing that, because $g \rhd_\ell f$, $(\Pi^N_\ell \cap \widehat{F}^g_\ell) \subseteq (\Pi^N_\ell \cap \widehat{F}^f_\ell)$, we obtain that

$$\theta^f_\ell(\Pi^N_\ell) = \theta^f_\ell(\Pi^N_\ell \cap \widehat{F}^f_\ell) \leq \theta^g_\ell(\Pi^N_\ell \cap \widehat{F}^g_\ell) < |\Pi^N_\ell \cap \widehat{F}^g_\ell| + 1 \leq |\Pi^N_\ell \cap \widehat{F}^f_\ell| + 1.$$

If follows that ℓ rejects f in Round N, a contradiction. □

It therefore remains only to prove Lemma A.1 in order to complete the proof of Theorem 6.

Proof of Lemma A.1. We prove each case in turn.

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

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

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

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

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

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

$$\min_{g \in \widehat{F}^f_\ell \cup \{f\}} \left\{ \widetilde{\theta}^g_\ell(G) \right\} \le \min_{g \in \widehat{F}^f_\ell \cup \{f\}} \left\{ \widetilde{\theta}^g_\ell(H) \right\} \le \min_{g \in \widehat{F}^f_\ell \cup \{f\}} \left\{ \widetilde{\theta}^g_\ell(G) + |\widehat{H}^g_\ell| - |\widehat{G}^g_\ell| \right\}. \tag{13}$$

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

$$\min_{g \in \widehat{F}^f_\ell \cup \{f\}} \left\{ \widetilde{\theta}^g_\ell(G) + |\widehat{H}^g_\ell| - |\widehat{G}^g_\ell| \right\} \le \widetilde{\theta}^h_\ell(G) + |\widehat{H}^h_\ell| - |\widehat{G}^h_\ell|. \tag{14}$$

In addition, we have that

$$|\widehat{H}^h_\ell| - |\widehat{G}^h_\ell| = |\widehat{H}^h_\ell \setminus \widehat{G}^h_\ell| \le |\widehat{H}^f_\ell \setminus \widehat{G}^f_\ell| = |\widehat{H}^f_\ell| - |\widehat{G}^f_\ell|, \tag{15}$$

where the two equalities follow from the fact that $\widehat{G}^h_\ell \subseteq \widehat{H}^h_\ell$ and $\widehat{G}^f_\ell \subseteq \widehat{H}^f_\ell$ while the fact that $(\widehat{H}^h_\ell \setminus \widehat{G}^h_\ell) \subseteq (\widehat{H}^f_\ell \setminus \widehat{G}^f_\ell)$ (as $h \in \widehat{F}^f_\ell$) implies the inequality. Combining (13), (14), and (15), we obtain

$$\min_{g \in \widehat{F}^f_\ell \cup \{f\}} \left\{ \widetilde{\theta}^g_\ell(G) \right\} \le \min_{g \in \widehat{F}^f_\ell \cup \{f\}} \left\{ \widetilde{\theta}^g_\ell(H) \right\} \le \widetilde{\theta}^h_\ell(G) + |\widehat{H}^f_\ell| - |\widehat{G}^f_\ell|. \tag{16}$$

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

$$\theta^f_\ell(G) = \min_{g \in \widehat{F}^f_\ell \cup \{f\}} \left\{ \widetilde{\theta}^g_\ell(G) \right\} \widetilde{\theta}^h_\ell(G) \quad \text{and} \quad \theta^f_\ell(H) = \min_{g \in \widehat{F}^f_\ell \cup \{f\}} \left\{ \widetilde{\theta}^g_\ell(H) \right\}. \tag{17}$$

Combining (16) and (17) yields

$$\theta^f_\ell(G) \le \theta^f_\ell(H) \le \theta^f_\ell(G) + |\widehat{H}^f_\ell| - |\widehat{G}^f_\ell|,$$

as required.

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

$$\theta^f_\ell(G) = \min_{h \in \widehat{G}^f_\ell} \left\{ \widetilde{\theta}^h_\ell \right\} \quad \text{and} \quad \theta^g_\ell(G) = \min_{h \in \widehat{G}^g_\ell} \left\{ \widetilde{\theta}^h_\ell \right\}.$$

Then, the fact that $g \rhd_\ell f$ implies $\widehat{G}^g_\ell \subseteq \widehat{G}^f_\ell$; therefore

$$\theta^g_\ell(G) = \min_{h \in \widehat{G}^g_\ell} \left\{ \widetilde{\theta}^h_\ell \right\} \ge \min_{h \in \widehat{G}^f_\ell} \left\{ \widetilde{\theta}^h_\ell \right\} = \theta^f_\ell(G),$$

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 $g \rhd_\ell f$ 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

$$\theta_\ell^g(G) - |\widehat{G}_\ell^g| > \theta_\ell^f(G) - |\widehat{G}_\ell^f|. \tag{18}$$

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 (18) 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 of the Lemma A.1. $\square$

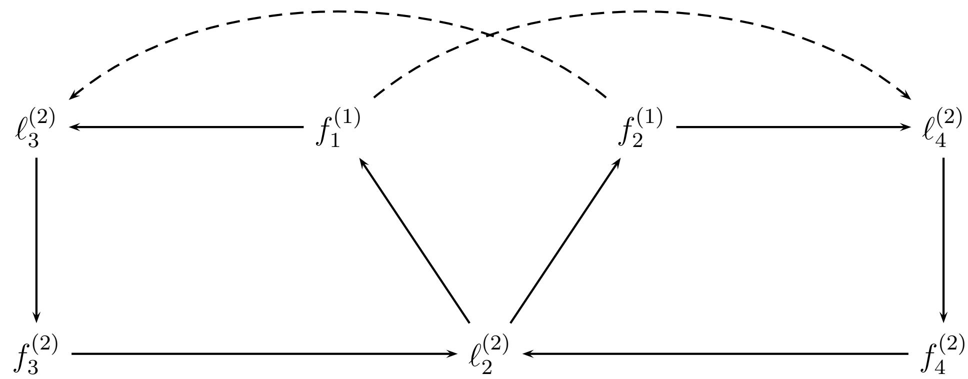

Figure B.1. Counterexample for the proof of Proposition B.1. $f \rightarrow \ell$: ℓ is f's first choice. $f \dashrightarrow \ell$: ℓ is f's second choice and f prefers ℓ to its endowment. $\ell \rightarrow f$: ℓ is f's endowment. Superscripts denote sizes and capacities, respectively.

Appendix B. Pareto Improvement from a Pareto-inefficient Endowment

Proposition B.1. *There is no strategy-proof mechanism that Pareto improves upon every endowment that is not Pareto-efficient.*

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

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

The preferences of families are as follows:

$$\succ_{f_1}: \ell_3, \ell_4, \boldsymbol{\ell_2}, \ldots \quad \succ_{f_2}: \ell_4, \ell_3, \boldsymbol{\ell_2}, \ldots \quad \succ_{f_3}: \ell_2, \boldsymbol{\ell_3}, \ldots \quad \succ_{f_4}: \ell_2, \boldsymbol{\ell_4}, \ldots,$$

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

$$\nu = \begin{array}{cc} & \begin{array}{cccc} f_1 & f_2 & f_3 & f_4 \end{array} \\ d_1 & \begin{pmatrix} 1 & 1 & 2 & 2 \end{pmatrix} \end{array} \qquad \kappa = \begin{array}{cc} & \begin{array}{ccc} \ell_2 & \ell_3 & \ell_4 \end{array} \\ d_1 & \begin{pmatrix} 2 & 2 & 2 \end{pmatrix} \end{array}.$$

For ease of exposition, we illustrate this counterexample in Figure B.1.

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 makes 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 $\succ'_{f_1}$: $\ell_3, \boldsymbol{\ell_2}, \ldots$. 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 $\succ'_{f_2}$: $\ell_4, \boldsymbol{\ell_2}, \ldots$ 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. $\square$

Appendix C. Envy-freeness and Stability

In this online appendix, we expand on Section 5.1 and detail how interference-freeness relates to existing solution concepts for respecting priorities: envy-freeness and stability. We discuss the logical relationship among the three solution concepts in Online Appendix C.1; the relationship between interference-freeness and envy-freeness in the special case where there is only one dimension in Online Appendix C.2; and the existence of stable matchings in Online Appendix C.3. Online Appendix C.4 contains the proofs of the results presented.

C.1. **Logical relationships between solution concepts.** We begin by formally defining envy-freeness (in the sense of elimination of justified envy).

Definition C.1. 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., $f \rhd_{\mu(f')} f'$.

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

Like interference-free matchings, envy-free matchings respect priorities but may be wasteful. However, the criterion is more extreme because it precludes *any* priority violation: if a family prefers a locality to its own, then only lower-priority family may be matched to that locality. In contrast, interference-freeness allows priority violations that are innocuous in the sense that the envying family does not have a claim to any unit used by the envied family. Therefore, if a family misses out on a more preferred locality, lower-priority families are not to blame. Therefore, interference-freeness still constitutes a strong criterion when it comes to respecting priorities; however, the relaxation it provides allows making families better off and reducing waste. We illustrate with an example.

Example C.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 $f_1 \rhd_{\ell_1} f_2 \rhd_{\ell_1} f_3$. There are two dimensions and the sizes and capacities are displayed below:

$$\nu = \begin{array}{c} \\ d_1 \\ d_2 \end{array} \begin{array}{c} \begin{array}{ccc} f_1 & f_2 & f_3 \end{array} \\ \begin{pmatrix} 1 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \end{array} \qquad \kappa = \begin{array}{c} \\ d_1 \\ d_2 \end{array} \begin{array}{c} \ell_1 \\ \begin{pmatrix} 1 \\ 1 \end{pmatrix} \end{array}.$$

The most efficient envy-free matching assigns f_1 to ℓ_1 and leaves the other two families unmatched.[20] 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. It is f_1 that prevents matching f_2 to ℓ_1, not f_3. Therefore, the matching in which

[20] If f_1 is unmatched, matching one or both of the other families to ℓ_1 would violate f_1's priority; therefore, the only other envy-free matching leaves all three families unmatched.

both f_1 and f_3 are matched to ℓ_1 appears to be a natural solution: it is more efficient and has no meaningful priority violations. That matching is interference-free since ℓ_1 can weakly accommodate f_3 alongside f_1 and f_2; in other words, the unit required by f_3 is claimed by neither f_1 nor f_2.

As interference-freeness allows for (innocuous) priority violation and envy-freeness does not, it is natural to think of interference-freeness as a relaxation of envy-freeness. Our first result formalizes this intuition.

Proposition C.1. Every envy-free matching is interference-free.

We next turn to stability and begin by formally defining it in our model.

Definition C.2. A matching μ is *(pairwise) stable* if there are 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* under 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.[21] If $|D| = 1$ and $\nu_{d_1}^f = 1$ for all $f \in F$, Definition C.2 collapses to the "elimination of justified envy" used in school choice and other object allocation settings (Abdulkadiroğlu and Sönmez, 2003).

Envy-freeness and stability are logically independent: Envy-freeness allows waste, which stability precludes; however, stability allows for some waste-eliminating priority violations. By Proposition C.1, it follows that stability does not imply interference-freeness. As interference-freeness allows waste but stability does not, we conclude that interference-freeness and stability are also logically independent. 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.[22] Perhaps surprisingly, although interference-freeness allows for some priority violations, we can establish an analogous relationship among stability, non-wastefulness, and interference-freeness.

Proposition C.2. If a matching is non-wasteful and interference-free, then it is stable.

[21] 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)).

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

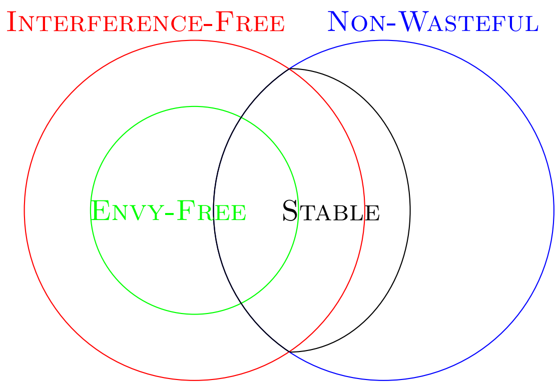

FIGURE C.2. Logical relationships among our solution concepts.

Proposition C.2, combined with the possible nonexistence of a stable matching (see Online Appendix C.3), formally establishes that an interference-free and non-wasteful matching may not exist. Figure C.2 summarizes the relationships among the solution concepts.

We present a small example to illustrate how stability and interference-freeness differ.

Example C.2. There are three families f_1, f_2, and f_3, one locality ℓ such that $f_1 \rhd_\ell f_2 \rhd_\ell f_3$, and one dimension d such that $\nu_d^{f_1} = 1$, $\nu_d^{f_2} = 2$, $\nu_d^{f_3} = 1$, and $\kappa_d^\ell = 2$.

The family-optimal interference-free matching assigns f_1 to ℓ and the other two families to the null. That matching is not stable since f_3 would prefer to be matched to ℓ and one unit of ℓ is not used. In fact, the unique stable matching assigns both f_1 and f_3 to ℓ (and f_2 to the null), but f_3 interferes with this matching so it is not interference-free. The example illustrates that a family can interfere even when higher-priority families could not make use of the capacity it uses: f_3 interferes with the unique stable matching because the unit it uses can be claimed by f_2, even though f_2 could not actually use that unit since ℓ cannot accommodate f_2 alongside f_1.

C.2. **Relationship between interference-freeness and envy-freeness with one dimension.** Example C.1 shows that the family-optimal interference-free matching—which assigns f_1 and f_3 to ℓ_1—may not be envy-free, and may therefore strictly dominate the family-optimal envy-free matching—which only assigns f_1 to ℓ_1. However, because Example C.1 has two dimensions, it leaves open the question of whether there exists a tighter relationship when $|D| = 1$. The following example shows that interference-freeness constitutes a relaxation of envy-freeness even when there is only one dimension.

Example C.3. There are two families f_1 and f_2, one locality ℓ such that $f_1 \rhd_\ell f_2$, and one dimension d such that $\nu_d^{f_1} = \nu_d^{f_2} = 1$ and $\kappa_d^\ell = 2$. There are four possible matchings:

$$\mu_1 = \begin{pmatrix} f_1 & f_2 \\ \ell_1 & \ell_2 \end{pmatrix} \quad \mu_2 = \begin{pmatrix} f_1 & f_2 \\ \ell_1 & \emptyset \end{pmatrix} \quad \mu_3 = \begin{pmatrix} f_1 & f_2 \\ \emptyset & \ell_2 \end{pmatrix} \quad \mu_4 = \begin{pmatrix} f_1 & f_2 \\ \emptyset & \emptyset \end{pmatrix}.$$

Matchings μ_1, μ_2, and μ_4 are envy-free but matching μ_3 is not because f_1 envies f_2. However, μ_3 is interference-free because ℓ can accommodate both f_1 and f_2; hence f_1 does not have a claim to the unit of capacity used by f_2.

In Example C.3, the family-optimal envy-free and interference-free matchings are identical and match both f_1 and f_2 to ℓ. As our next result formalizes, this is not a coincidence.

Proposition C.3. If $|D| = 1$, then the KDA and TKDA mechanisms are envy-free.

As the KDA mechanism outputs the family-optimal interference-free matching (Theorem 5) and every interference-free matching is envy-free (Proposition C.1), Proposition C.3 implies the following corollary.

Corollary C.1. If $|D| = 1$, the family-optimal envy-free and interference-free matchings coincide.

Taken together, our results show that interference-freeness constitutes a relaxation of envy-freeness even when there is only one dimension; however, that relaxation does not have practical consequences since in that case, both our mechanisms output envy-free matchings. In contrast, using interference-freeness rather than envy-freeness matters when there are multiple dimensions. In Example C.1, both the KDA and TKDA algorithm match f_1 and f_3 to ℓ_1, which envy-freeness would preclude. More generally, our simulation results suggest that the efficiency gains associated with using interference-freeness instead of envy-freeness as a criterion for respecting priorities can be large (see Figure F.2).

C.3. **(Non-)existence of stable matchings.** While stable matchings always exist in school choice models, they do not exist in ours (even when $|D| = 1$). In fact, determining whether a stable matching exists in our model (even when $|D| = 1$) is a computationally intractable problem (McDermid and Manlove, 2010). We first reproduce an example from McDermid and Manlove (2010) in which no stable matching exists and then provide a condition on the priority profile for the existence of a stable matching.

Example C.4 (McDermid and Manlove (2010))**.** There are three families, two localities, and one dimension. The family sizes and locality capacities are

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

The preferences and priorities are:

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

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_3 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 why stable matchings might not exist is that the choice function induced by ℓ_1's priorities is not substitutable (Roth, 1984) 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.

We next provide a condition on the priority profile for the existence of a stable matching.

Definition C.3. A priority profile $\rhd$ is *aligned* if for any $f, g \in F$ such that $\nu^f \neq \nu^g$ and any $\ell, \ell' \in L \setminus \{\emptyset\}$, $f \rhd_\ell g$ if and only if $f \rhd_{\ell'} g$.

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 C.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 $\{F_1, F_2, \ldots, F_n\}$ such that for any $i = 1, \ldots, n$, all families in F_i have the same size and for all $j < i$, all families in F_j have a higher priority at 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 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 C.4. As in school choice, stable matchings under size monotonicty and priority alignment can be found in polynomial time in our model.

C.4. **Proofs of results in Online Appendix C.**

Proof of Proposition C.1. We prove the contrapositive: Consider a matching μ that is not interference-free, we need to show that μ is not envy-free. By definition, there exists a family

f such that $\mu(f)$ cannot weakly accommodate f alongside the set

$$G = \{g \in F : g \rhd_{\mu(f)} f \text{ and } \mu(f) \succeq_g \mu(g)\}.$$

It follows that there exists a dimension d such that $\nu_d^f + \sum_{g \in G} \nu_d^g > \kappa_d^{\mu(f)}$. If $G \setminus \mu(f) = \varnothing$, then $\mu(f)$ cannot accommodate all families matched to it, a contradiction. Therefore, there exists a family $f' \in G \setminus \mu(f)$ and by definition $f' \rhd_{\mu(f)} f$ and $\mu(f) \succ_{f'} \mu(f')$. We conclude that f' envies f, so μ is not envy-free. □

Proof of Proposition C.2. We prove the contrapositive. Let μ be a matching that is non-wasteful and not stable. We need to show that μ is not interference-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 $d \in D$,

$$\nu_d^f + \sum_{g \in (\widehat{F}_\ell^f \cap \mu(\ell))} \nu_d^g \leq \kappa_d^\ell. \tag{19}$$

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

$$\nu_{d'}^f + \sum_{g \in \mu(\ell)} \nu_{d'}^g > \kappa_{d'}^\ell. \tag{20}$$

Consider now the families in $\mu(\ell) \setminus \widehat{F}_\ell^f$ whose size in dimension d' is at least 1. Inequalities (19) and (20) 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 1 (equivalently, f' is the lowest-priority family in $\mu(\ell)$ such that $\nu_{d'}^{f'} > 0$). We show that f' interferes with μ; that is, we show that ℓ cannot weakly accommodate f' alongside $\widehat{F}_\mu^{f'} = \{g \in F : g \rhd_\ell f' \text{ and } \ell \succeq_g \mu(g)\}$. By the definition of f', $\nu_{d'}^{f'} + \sum_{g \in (\widehat{F}_\ell^{f'} \cap \mu(\ell))} \nu_{d'}^g = \sum_{g \in \mu(\ell)} \nu_{d'}^g$, which combined with (20) implies that

$$\nu_{d'}^f + \nu_{d'}^{f'} + \sum_{g \in (\widehat{F}_\ell^{f'} \cap \mu(\ell))} \nu_{d'}^g = \nu_{d'}^f + \sum_{g \in \mu(\ell)} \nu_{d'}^g > \kappa_{d'}^\ell. \tag{21}$$

Inequality (21) 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.,

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

so ℓ cannot weakly accommodate f' alongside $\widehat{F}_\mu^{f'}$ and μ is not interference-free. □

Proof of Proposition C.3. 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 $i = 1, \ldots, N$, $\widehat{R}^i_f(\ell)$ 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 $f \rhd_{\ell'} f'$. By construction, there exists a Round i in which ℓ' rejects f; hence ℓ' cannot weakly accommodate f alongside $\widehat{R}^i_f(\ell')$. Again by construction, we have that $\widehat{R}^i_f(\ell') \subseteq \widehat{R}^N_f(\ell')$ so we conclude that ℓ' cannot weakly accommodate f alongside $\widehat{R}^N_f(\ell')$.

As d is the only dimension, it must be that $\nu^f_d > 0$, and therefore the fact that ℓ' cannot weakly accommodate f alongside $\widehat{R}^N_f(\ell')$ implies by definition that $\nu^f_d + \sum_{g \in \widehat{R}^N_f(\ell')} \nu^g_d > \kappa^{\ell'}_d$. As $f \rhd f'$, it follows that $\nu^{f'}_d + \sum_{g \in \widehat{R}^N_{f'}(\ell')} \nu^g_d > \kappa^{\ell'}_d$. Therefore, ℓ' cannot weakly accommodate f' alongside $\widehat{R}^N_{f'}(\ell')$ 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 $i = 1, \ldots, N$ and any locality $\ell \in L$, Π^i_ℓ 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 $G \subseteq F$, $\theta^f_\ell(G)$ denotes f's threshold for ℓ (calculated by Algorithm 5) when the families in G propose to ℓ. Moreover, we will invoke Lemma A.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 $f \rhd_{\ell'} f'$, and that for any $\tilde{f} \neq f$ such that $\ell' \succ_{\tilde{f}} \mu(\tilde{f})$, $f \rhd_{\ell'} \tilde{f}$. As $\ell' \succ_f \mu(f)$, by construction there exists a Round i in which ℓ' permanently rejects f so we have that

$$\theta^f_{\ell'}(\Pi^i_{\ell'}) < |\Pi^i_{\ell'} \cap \widehat{F}^f_{\ell'}| + 1. \tag{22}$$

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

$$\theta^f_{\ell'}(\Pi^N_{\ell'}) < |\Pi^N_{\ell'} \cap \widehat{F}^f_{\ell'}| + 1. \tag{23}$$

If $\theta^f_{\ell'}(\Pi^N_{\ell'}) = \infty$, then by Lemma A.1(i), $\theta^f_{\ell'}(\Pi^i_{\ell'}) = \infty$, which contradicts (22) and, if $\theta^f_{\ell'}(\Pi^N_{\ell'}) = 0$, (23) holds trivially. Therefore, it remains to show that (23) holds in the case where $\theta^f_{\ell'}(\Pi^N_{\ell'}) \in \mathbb{Z}_{>0}$. 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 $\Pi^i_{\ell'} \cap \widehat{F}^f_{\ell'} \subseteq \Pi^N_{\ell'} \cap \widehat{F}^f_{\ell'}$. By Lemma A.1(ii), it follows that

$$\theta^f_{\ell'}(\Pi^N_{\ell'} \cap \widehat{F}^f_{\ell'}) \leq \theta^f_{\ell'}(\Pi^i_{\ell'} \cap \widehat{F}^f_{\ell'}) + |\Pi^N_{\ell'} \cap \widehat{F}^f_{\ell'}| - |\Pi^i_{\ell'} \cap \widehat{F}^f_{\ell'}|.$$

By construction (see Algorithm 5), f's threshold at ℓ' only depends on higher-priority families so $\theta^f_{\ell'}(\Pi^N_{\ell'} \cap \widehat{F}^f_{\ell'}) = \theta^f_{\ell'}(\Pi^N_{\ell'})$ and $\theta^f_{\ell'}(\Pi^i_{\ell'} \cap \widehat{F}^f_{\ell'}) = \theta^f_{\ell'}(\Pi^i_{\ell'})$; hence we have that

$$\theta^f_{\ell'}(\Pi^N_{\ell'}) \leq \theta^f_{\ell'}(\Pi^i_{\ell'}) + |\Pi^N_{\ell'} \cap \widehat{F}^f_{\ell'}| - |\Pi^i_{\ell'} \cap \widehat{F}^f_{\ell'}|,$$

which, combined with (22), implies (23).

Having established that inequality (23) holds, we now use it to reach the contradiction. We consider separately two cases: $\theta^{f'}_{\ell'}(\Pi^N_{\ell'}) \neq \infty$ and $\theta^{f'}_{\ell'}(\Pi^N_{\ell'}) = \infty$.

Case 1: $\theta^{f'}_{\ell'}(\Pi^N_{\ell'}) \neq \infty$. As $f \rhd_{\ell'} f'$, we have that $\Pi^N_{\ell'} \cap \widehat{F}^f_{\ell'} \subseteq \Pi^N_{\ell'} \cap \widehat{F}^{f'}_{\ell'}$ and, by Lemma A.1(iii), $\theta^{f'}_{\ell'}(\Pi^N_{\ell'}) \leq \theta^f_{\ell'}(\Pi^N_{\ell'})$. Then, (23) implies that $\theta^{f'}_{\ell'}(\Pi^N_{\ell'}) < |\Pi^N_{\ell'} \cap \widehat{F}^{f'}_{\ell'}| + 1$. 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: $\theta^{f'}_{\ell'}(\Pi^N_{\ell'}) = \infty$. By definition (see Algorithm 5), the case assumption implies that ℓ' can weakly accommodate f' alongside $\widehat{F}^{f'}_{\ell'}$. As d is the only dimension, we have that $\nu^{f'}_d > 0$ so, by definition: $\nu^{f'}_d + \sum_{g \in \widehat{F}^{f'}_{\ell'}} \nu^g_d \leq \kappa^{\ell'}_d$. It follows that $\nu^f_d + \sum_{g \in \widehat{F}^f_{\ell'}} \nu^g_d < \kappa^{\ell'}_d$; hence ℓ' can (weakly) accommodate f alongside $\widehat{F}^f_{\ell'}$. By definition (see Algorithm 5), we conclude that $\theta^f_{\ell'}(\Pi^N_{\ell'}) = \infty$, which contradicts (23). □

Proof of Proposition C.4. Suppose that sizes are monotonic and the priority profile is aligned. We introduce the Sequential Deferred Acceptance (SDA) algorithm (Algorithm C.1) 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.

We first show that the SDA algorithm ends after finitely many rounds. If there is a directed cycle $(f_1, f_2, \ldots, f_n)$ in graph $\mathbb{G}^1$, then, for every $\ell \in L \setminus \{\emptyset\}$, $f_1 \rhd_\ell f_2 \rhd_\ell \ldots \rhd_\ell f_n \rhd_\ell f_1$, a contradiction. Therefore, $\mathbb{G}^1$ is a directed acyclic graph. By construction, for all $i > 1$, $\mathbb{G}^i$ is a directed acyclic graph 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 $\widetilde{F}^i$ at which no family is pointing is nonempty. It follows that at least one family gets permanently matched in each round; hence the algorithm ends after at most $|F|$ rounds.

Algorithm C.1. Sequential Deferred Acceptance

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

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

Round $i \geq 1$

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

Permanently match the families in $\widetilde{F}^i$ to the localities using the family-proposing Deferred Acceptance algorithm setting the capacity of each 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 $\widetilde{F}^i$ and all edges adjacent to them. For each locality $\ell \in L$, let $\widetilde{F}_\ell^i$ be the set of families that have been permanently matched to ℓ in Round i. Update the counter of each locality ℓ as follows: $c_\ell^{i+1} = c_\ell^i - \sum_{f \in \widetilde{F}_\ell^i} \nu^f$. Continue to Round $i + 1$.

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

Let Round i be the round in which family f is permanently matched to $\mu^{\text{SDA}}(f)$, i.e., $f \in \widetilde{F}^i$. Consider the Round i "submarket," in which the DA algorithm permanently matches the families in $\widetilde{F}^i$ 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 $f' \in \widetilde{F}^i$. (If this were not the case, then, as the priority profile is aligned, one of f or f' would have a higher priority at all localities and would point at the other, a contradiction.) 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 $\ell \succ_f \mu^{\text{SDA}}(f)$ 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 $\widehat{F}_\ell^f \cap \widetilde{F}_\ell^i$. 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 $\cup_{j=1}^i \widetilde{F}_\ell^j \subseteq \mu^{\text{SDA}}(\ell)$. We therefore conclude that ℓ cannot accommodate f alongside $\widehat{F}_\ell^f \cap \mu^{\text{SDA}}(\ell)$, as required. $\square$

Appendix D. Efficiency of the TKDA Algorithm

In Online Appendix D.1, we show that the TKDA algorithm is more efficient than any envy-free mechanism that satisfies the cardinal monotonicity condition. A similar result replacing envy-freeness interference-freeness holds when $|D| = 1$. In Online Appendix D.2, we provide a lower bound for the efficiency of the KDA and TKDA algorithms. We introduce a solution to improve the efficiency of the TKDA algorithm in Online Appendix D.3 and show in Online Appendix D.4 that it cannot be generalized. All proofs are in Online Appendix D.5.

D.1. **TKDA and cardinal monotonicity.** As noted in Section 5.2, a key condition for a deferred-acceptance-type procedure to be strategy-proof is that its associated choice function satisfies the cardinal monotonicity condition. In this section, we show that the choice function of the TKDA algorithm tentatively accepts every family that is tentatively accepted by any choice function satisfying envy-freeness and cardinal monotonicity conditions. When $|D| = 1$, we show that the choice function of the TKDA algorithm maximizes the number of tentatively accepted families among choice functions that satisfy the interference-freeness and cardinal monotonicity conditions. When $|D| > 1$, we argue that the problem of maximizing the number of tentatively accepted families subject to the interference-freeness and cardinal monotonicity conditions is computationally intractable.

Fixing a locality ℓ, recall from the proof of Theorem 6 that a choice function $C_\ell : F \to F$ determines, for any subset $G \subseteq F$ of proposition families, which families the locality tentatively accepts ($C_\ell(G)$) or permanently rejects ($G \backslash C_\ell(G)$). We now define four properties of a choice function:

(1) C_ℓ is *feasible* if, for every $G \subseteq F$, ℓ can accommodate $C_\ell(G)$;
(2) C_ℓ is *envy-free* if, for every $G \subseteq F$, every family in $C_\ell(G)$ has a higher priority than any family in $G \setminus C_\ell(G)$;
(3) C_ℓ is *interference-free* if, for every $G \subseteq F$, ℓ can weakly accommodate any family $f \in C_\ell(G)$ alongside $G \cap \widehat{F}_\ell^f$;
(4) C_ℓ *satisfies the cardinal monotonicity condition* if $|C_\ell(G)| \leq |C_\ell(H)|$ for every $G \subseteq H \subseteq F$.

Denote by C_ℓ^T the choice function of any locality ℓ in the TKDA algorithm: For each family, a threshold is calculated by Algorithm 5 and the family is tentatively accepted if its threshold does not exceed its priority rank among proposing families, otherwise the family is permanently rejected. We first assess that choice function against the four properties previously defined.

Proposition D.1. The choice function C_ℓ^T is feasible, interference-free, and satisfies the cardinal monotonicity condition. Moreover, if $|D| = 1$, then C_ℓ^T is envy-free.

Proposition D.1 is closely related to the properties of the TKDA mechanism (Theorem 6 and Proposition C.3). We now compare C_ℓ^T to other choice functions that have similar properties.

Proposition D.2. For any feasible and envy-free choice function C_ℓ that satisfies the cardinal monotonicity condition, and for any subset of families $G \subseteq F$, we have that $C_\ell(G) \subseteq C_\ell^T(G)$.

Proposition D.2 provides a lower bound for the efficiency of the TKDA mechanism as its choice function is at least as efficient as any envy-free choice function that satisfies the cardinal monotonicity condition: any family that is accepted by such a choice function is accepted by the TKDA choice function. Moreover, the TKDA mechanism can do even better by allowing families that can be weakly accommodated alongside all higher-priority families to be tentatively accepted even when some higher-priority families are permanently rejected.

However, Proposition D.2 leaves open the question of what can be said about feasible and interference-free choice functions that satisfy the cardinal monotonicity condition. Our next result provides an answer for the case where there is only one dimension.

Proposition D.3. If $|D| = 1$, then for any feasible and interference-free choice function C_ℓ that satisfies the cardinal monotonicity condition, and for any subset of families $G \subseteq F$, we have that $|C_\ell(G)| \leq |C_\ell^T(G)|$. Moreover, every family in $C_\ell^T(G) \setminus C_\ell(G)$ has a higher priority than any family in $C_\ell(G) \setminus C_\ell^T(G)$.

When $|D| = 1$, the TKDA choice function is envy-free (Proposition D.1) and accepts the highest-priority families until accepting the next family would violate either the locality's capacity or the cardinal monotonicity condition. Other feasible interference-free choice functions satisfying the cardinal monotonicity condition might not accept a larger number of families but might accept lower-priority families. The next example illustrates this.

Example D.1. There are four families f_1, f_2, f_3, and f_4 and one locality ℓ. The priority list of ℓ is $f_1 \rhd_\ell f_2 \rhd_\ell f_3 \rhd_\ell f_4$. There is one dimension and the sizes and capacities are displayed below:

$$\nu = \begin{array}{c c} & \begin{array}{cccc} f_1 & f_2 & f_3 & f_4 \end{array} \\ d_1 & \begin{pmatrix} 2 & 1 & 1 & 1 \end{pmatrix} \end{array} \qquad \kappa = \begin{array}{c c} & \ell \\ d_1 & \begin{pmatrix} 3 \end{pmatrix} \end{array}.$$

One can verify that $C_\ell^T(\{f_2, f_3, f_4\}) = \{f_2, f_3\}$: the rank of f_2 is ∞ while the rank of both f_3 and f_4 is 2. Therefore, f_4 is rejected even though ℓ can accommodate all three families. This occurs because of the cardinal monotonicity condition: when all four families propose, only f_1 and f_2 can be accepted; therefore, the cardinal monotonicity condition dictates that at most two families be accepted when f_2, f_3, and f_4 propose. Then, any feasible and interference-free choice function that satisfies the cardinal monotonicity condition can accept at most two families, i.e., no more than the TKDA choice function. However, not every such choice function needs to accept the same two families: f_4 could be accepted instead of either f_2 or f_3. □

Proposition D.3 provides a strong efficiency result for the TKDA mechanism when there is only one dimension: its choice function cannot accept more families without violating the interference-freeness or cardinal monotonicity property. When there are multiple dimensions, Proposition D.2 continues to provide a lower bound for the efficiency of the TKDA mechanism but Proposition D.3 breaks down: there may be feasible and interference-free choice functions that satisfy the cardinal monotonicity condition and accept more families than C_ℓ^T.

Example D.2. There are four families f_1, f_2, f_3, and f_4 and one locality ℓ. The priority list of ℓ is $f_1 \rhd_\ell f_2 \rhd_\ell f_3 \rhd_\ell f_4$. There are two dimensions and the sizes and capacities are displayed below:

$$\nu = \begin{array}{c c} & \begin{array}{cccc} f_1 & f_2 & f_3 & f_4 \end{array} \\ \begin{array}{c} d_1 \\ d_2 \end{array} & \begin{pmatrix} 1 & 0 & 1 & 1 \\ 0 & 1 & 1 & 0 \end{pmatrix} \end{array} \qquad \kappa = \begin{array}{c c} & \begin{array}{c} \ell \end{array} \\ \begin{array}{c} d_1 \\ d_2 \end{array} & \begin{pmatrix} 2 \\ 1 \end{pmatrix} \end{array}.$$

Locality ℓ cannot accommodate $\{f_2, f_3\}$; therefore, when f_3 and f_4 propose, f_3 and f_4's thresholds are 1 so $C_\ell^T(\{f_3, f_4\}) = \{f_3\}$. When f_2 proposes alongside f_3 and f_4, f_3 and f_4's thresholds are 0 so $C_\ell^T(\{f_2, f_3, f_4\}) = \{f_2\}$. Consider the alternative choice function C_ℓ such that $C_\ell(\{f_3, f_4\}) = \{f_3, f_4\}$, $C_\ell(\{f_2, f_3, f_4\}) = \{f_2, f_4\}$, and $C_\ell(G) = C_\ell^T(G)$ for every other subset of families G. First, C_ℓ is feasible since ℓ can accommodate $\{f_3, f_4\}$ and $\{f_2, f_4\}$. Second, C_ℓ is interference-free as $\nu_{d_2}^{f_4} = 0 < \nu_{d_2}^{f_3}$ and ℓ can weakly accommodate f_4 alongside $\{f_2, f_3\}$ (hence alongside $\{f_3\}$). Third, C_ℓ satisfies the cardinal monotonicity condition since $|C_\ell(\{f_1, f_3, f_4\})| = |\{f_1, f_3\}| = 2$ and $|C_\ell(\{f_1, f_2, f_3, f_4\})| = |\{f_1, f_2\}| = 2$. □

Therefore, Example D.2 shows that Proposition D.3 does not hold when there are multiple dimensions. Unfortunately, as the next example shows, improving the efficiency of the TKDA choice function is computationally difficult.

Example D.3. There are three families f_{12}, f_3, and f_4 and one locality ℓ. The priority list of ℓ is $f_{12} \rhd_\ell f_3 \rhd_\ell f_4$. There are two dimensions and the sizes and capacities are displayed below:

$$\nu = \begin{array}{c c} & \begin{array}{ccc} f_{12} & f_3 & f_4 \end{array} \\ \begin{array}{c} d_1 \\ d_2 \end{array} & \begin{pmatrix} 1 & 1 & 1 \\ 1 & 1 & 0 \end{pmatrix} \end{array} \qquad \kappa = \begin{array}{c c} & \begin{array}{c} \ell \end{array} \\ \begin{array}{c} d_1 \\ d_2 \end{array} & \begin{pmatrix} 2 \\ 1 \end{pmatrix} \end{array}.$$

This example is identical to Example D.2 except that the families f_1 and f_2 are merged into family f_{12}. For any feasible and interference-free choice function C_ℓ that satisfies the cardinal monotonicity condition, $C_\ell(\{f_{12}, f_3, f_4\}) \subseteq \{f_{12}\}$ as ℓ can weakly accommodate neither f_3 alongside f_{12} nor f_4 alongside $\{f_{12}, f_3\}$. Therefore, in contrast to Example D.2, if

f_3 and f_4 propose to ℓ the cardinal monotonicity condition precludes ℓ from accepting both families. This difference occurs even though the sizes of f_3 and f_4 as well as the number of units required by higher-priority families have not changed. □

The TKDA algorithm is computationally simple because it relies exclusively on the aggregate sizes of the higher-priority families to determine a family's threshold. Making the choice function more efficient requires considering the sizes of all higher-priority families and checking what occurs when any *subset* of them proposes, which is computationally challenging as the number of subsets grows exponentially with the number of families.

For a small enough problem, such a choice function can be constructed as follow. Start with the choice function C_ℓ^K of the KDA algorithm, which tentatively accepts each family f if ℓ can weakly accommodate f alongside all higher-priority families and permanently rejects f otherwise. That choice function is feasible and interference-free but may not satisfy the cardinal monotonicity condition. Let $C_\ell^M(G) = \{f \in G \mid |\widehat{G}_\ell^f| < \min_{H \supseteq G}\{|C_\ell^K(H)|\}\}$. If C_ℓ^K satisfies the cardinal monotonicity condition, then C_ℓ^M and C_ℓ^K are identical. Otherwise, every time the cardinal monotonicity condition is violated, i.e., for every two subsets of families $G \subseteq H \subseteq F$ such that $|C_\ell^K(G)| > |C_\ell^K(H)|$, the lowest-priority families in $C_\ell^K(G)$ are permanently rejected, until the violation is resolved. C_ℓ^M is efficient as it maximizes the number of families that are tentatively accepted subject to satisfying the cardinal monotonicity condition; however, computing C_ℓ^M is demanding as it requires looking at each subset of families. Essentially, the choice function C_ℓ^T in our TKDA algorithm is a greedy version of C_ℓ^M although C_ℓ^M and C_ℓ^T are equivalent when $|D| = 1$ (by Proposition D.3).

Finally, our analysis leaves open the question of what strategy-proof and interference-free mechanisms can be constructed beyond the cardinal monotonicity condition. The cardinal monotonicity condition is sufficient for a deferred-acceptance-type procedure to be strategy-proof. Hatfield and Milgrom (2005) showed that if the cardinal monotonicity condition is not satisfied, then there are instances in which a deferred-acceptance-type procedure is not strategy-proof. However, our TKDA with Clinches (TKDAC), which we introduce in Online Appendix D.3, shows that it is possible to identify violations of the cardinal monotonicity condition that do not affect strategy-proofness. We leave for future research the question of whether and how this can be done beyond clinches.

D.2. **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 D.4. 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 D.4 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 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.

D.3. **TKDA with Clinching.** 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 D.1) starts with a Clinching Round that creates a new priority profile. The TKDA algorithm is then run using the new priority profile.

Proposition D.5. The TKDAC mechanism is strategy-proof and interference-free. Moreover, $\mu^{\text{TKDAC}} \succeq \mu^{\text{TKDA}}$.

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 $\mu^{\text{TKDA}}(f) = \ell$: 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 each 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 at 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.

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 $\rhd_{\ell_1} : f_1, f_2, f_3, \ldots$. There is only one dimension d_1 and the sizes and capacities are

$$\begin{array}{cc} & \begin{array}{cccc} f_1 & f_2 & f_3 & \ell_1 \end{array} \\ d_1 & \left(\begin{array}{cccc} 2 & 1 & 1 & 2 \end{array} \right). \end{array}$$

Consider what happens if f_2 and f_3 both propose to ℓ_1 in some round of the TKDA algorithm. The thresholds are $\theta^{f_1}_{\ell_1} = \infty$ (since f_1 has the highest priority at ℓ_1 and ℓ_1 can accommodate

Algorithm D.1. TKDA with Clinching (TKDAC)

Clinching Round

Step 0

No locality rejects or proposes to any family and no family clinches any locality. Set $\widetilde{\rhd}^1 = \rhd$ and continue to Step 1.

Step $j \geq 1$

(a) Each 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 $\widetilde{\rhd}^j_\ell$.

(b) Each locality ℓ proposes to family f if ℓ can weakly accommodate f alongside all families that are higher than f on $\widetilde{\rhd}^j_\ell$.

(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 $\widetilde{\rhd} = \widetilde{\rhd}^j$ and continue to the TKDA algorithm.

(e) Construct $\widetilde{\rhd}^{j+1}$ as follows, then continue to Step $j+1$. For each $\ell \in L$ and each $f, f' \in F$ with $f \rhd_\ell f'$,

- $f' \widetilde{\rhd}^{j+1}_\ell f$ 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;
- $f \widetilde{\rhd}^{j+1}_\ell f'$ otherwise.

TKDA

Run the TKDA algorithm with the priority profile $\widetilde{\rhd}$.

f_1 on its own) and $\theta^{f_2}_{\ell_1} = \theta^{f_3}_{\ell_1} = 1$ (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 $1 = \theta^{f_2}_{\ell_1}$) and permanently rejects f_3's proposal (since f_3's priority rank among proposing families is $2 > 1 = \theta^{f_3}_{\ell_1}$). 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 $\widetilde{\rhd}$ such that $f_2 \widetilde{\rhd}_{\ell_1} f_3 \widetilde{\rhd}_{\ell_1} f_1$. If the TKDA algorithm is run with the new priority profile $\widetilde{\rhd}$ and f_2 and f_3 both propose to ℓ_1, we have that $\theta^{f_2}_{\ell_1} = \theta^{f_3}_{\ell_1} = \infty$ since f_2 and f_3 now have the two highest priorities at 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 in 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

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

Table D.1. TKDAC* algorithm applied to Example D.4.

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 E.3 provides an example of the TKDAC algorithm.

D.4. **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 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 D.1: 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* algorithm 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 D.4. There are three families, three localities, and one dimension. The preferences, priorities, family sizes, and locality capacities are

$$\succ_{f_1}:\ \ell_1,\ell_2,\ell_3 \qquad \succ_{f_2}:\ \ell_2,\ell_1,\ell_3 \qquad \succ_{f_3}:\ \ell_1,\ell_2,\ell_3$$
$$\rhd_{\ell_1}:\ f_2,f_3,f_1 \qquad \rhd_{\ell_2}:\ f_1,f_2,f_3 \qquad \rhd_{\ell_3}:\ f_1,f_2,f_3$$

$$\nu = \begin{array}{cc} & \begin{array}{ccc} f_1 & f_2 & f_3 \end{array} \\ d_1 & \begin{pmatrix} 2 & 1 & 1 \end{pmatrix} \end{array} \qquad \kappa = \begin{array}{cc} & \begin{array}{ccc} \ell_1 & \ell_2 & \ell_3 \end{array} \\ d_1 & \begin{pmatrix} 2 & 2 & 2 \end{pmatrix} \end{array}.$$

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 family's third- and last-preference locality, the modified Clinching Round does not affect the localities' priorities.

As a result, the priority profile $\widetilde{\rhd}$ that is used in the TKDA algorithm is the same as the original priority profile $\rhd$.[23]

The TKDA algorithm—summarized in Table D.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 $\widetilde{\rhd}_{\ell_2} : f_2, f_3, f_1$.[24] Families f_2 and f_3 clinch locality f_1, but this proposal is inconsequential since f_3 already has the highest priority at ℓ_3. Therefore, the modified Clinching Round does not modify the priorities of ℓ_1 and ℓ_3, i.e., $\widetilde{\rhd}_{\ell_1} = \rhd_{\ell_1}$ and $\widetilde{\rhd}_{\ell_3} = \rhd_{\ell_3}$.

The TKDA algorithm is then run with the priority profile $\widetilde{\rhd}$. 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 D.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).

D.5. **Proofs of results in Online Appendix D.** Throughout, for any family $f \in F$, any locality ℓ, and any subset of families $G \subseteq F$, we denote by $\widehat{G}_\ell^f = G \cap \widehat{F}_\ell^f$ the subset of families in G that have a higher priority for ℓ than f. We also denote by $\widetilde{\theta}_\ell^f(G)$ (resp. $\theta_\ell^f(G)$) the temporary threshold (resp. the threshold) of f for ℓ when all families in G propose to ℓ, as determined by Algorithm 5.

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

[24]Family 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.

Proof of Proposition D.1. We have already shown in the proof of Theorem 6 that C_ℓ^T is feasible and satisfies the cardinal monotonicity condition. It remains to show that C_ℓ^T is interference-free and, when $|D| = 1$, envy-free.

Interference-freeness. Consider a subset of families G and a family $f \in C_\ell^T(G)$. We need to show that ℓ can weakly accommodate f alongside $\widehat{G}_\ell^f$. Suppose towards a contradiction that this is not the case. Then, by definition (see Algorithm 5), $\theta_\ell^f(G) = 0$ and so $f \notin C_\ell^T(G)$, a contradiction.

Envy-freeness when $|D| = 1$. Suppose that $|D| = 1$ and consider a subset of families G as well as two families $f \in C_\ell^T(G)$ and $g \in G \setminus C_\ell^T(G)$. We need to show that $f \rhd_\ell g$. Suppose towards a contradiction that $g \rhd_\ell f$.

If $\theta_\ell^f(G) = \infty$, then by definition (see Algorithm 5), ℓ can weakly accommodate f alongside $\widehat{G}_\ell^f$. As $g \rhd_\ell f$ and $|D| = 1$, it follows that ℓ can weakly accommodate g alongside $\widehat{G}_\ell^g$ so $g \in C_\ell^T(G)$, a contradiction.

If $\theta_\ell^f(G) \neq \infty$, then by definition (see Algorithm 5) $\theta_\ell^f(G) \leq \theta_\ell^g(G)$ (since $g \rhd_\ell f$). As $f \in C_\ell^T(G)$ and $g \rhd_\ell f$, it follows that $\theta_\ell^g(G) \geq \theta_\ell^f(G) \geq |\widehat{G}_\ell^f| + 1 \geq |\widehat{G}_\ell^g| + 1$ so $g \in C_\ell^T(G)$, a contradiction. □

Proof of Proposition D.2. We need to show that $C_\ell(G) \subseteq C_\ell^T(G)$. The result holds trivially if $C_\ell^T(G) = G$; otherwise, let f be the highest-priority family in $G \setminus C_\ell^T(G)$. As C_ℓ is envy-free, it is sufficient to show that $f \notin C_\ell(G)$. If ℓ cannot accommodate $\widehat{G}_\ell^f \cup \{f\}$, then $f \notin C_\ell(G)$ since C_ℓ is feasible and envy-free. If ℓ can accommodate $\widehat{G}_\ell^f \cup \{f\}$, then ℓ can weakly accommodate f alongside $\widehat{G}_\ell^f$; hence, f's temporary threshold $\widetilde{\theta}_\ell^f(G)$ (as calculated by Algorithm 5) is at least $|\widehat{G}_\ell^f| + 1$. As f has the $(|\widehat{G}_\ell^f| + 1)$-st highest priority among the families in G and $f \notin C_\ell^T(G)$, f's threshold θ_ℓ^f is at most $|\widehat{G}_\ell^f|$. Therefore, we have that $\theta_\ell^f < \widetilde{\theta}_\ell^f(G)$ so there exists a family $g \in \widehat{F}_\ell^f$ such that $\widetilde{\theta}_\ell^g(G) \leq |\widehat{G}_\ell^f|$. Then, there exists a subset of families $H \subseteq \widehat{F}_\ell^g$ with $|H| = |\widehat{G}_\ell^f|$ and $\widehat{G}_\ell^g \subseteq H$ such that ℓ cannot weakly accommodate g alongside H, which implies that ℓ cannot accommodate $H \cup \{g\}$. Then, all families in $C_\ell(G \cup H \cup \{g\})$ have a higher priority than g at ℓ, as otherwise the assumption that C_ℓ is envy-free would imply that $(H \cup \{g\}) \subseteq C_\ell(G \cup H \cup \{g\})$, which would violate the assumption that C_ℓ is feasible. As $\widehat{G}_\ell^g \subseteq H$, it follows that $C_\ell(G \cup H \cup \{g\}) \subseteq H$, which combined with the fact that $|\widehat{G}_\ell^f| = |H|$ implies that $|C_\ell(G \cup H \cup \{g\})| \leq |\widehat{G}_\ell^f|$. Finally, as C_ℓ satisfies the cardinal monotonicity condition, we have that $|C_\ell(G)| \leq |\widehat{G}_\ell^f|$; hence, the assumption that C_ℓ is envy-free yields $f \notin C_\ell(G)$. □

Proof of Proposition D.3. Recall that, by assumption, there is only one dimension, which we denote throughout by d. Therefore, weak accommodation is equivalent to accommodation.

Denote by C_ℓ^K the choice function of the KDA algorithm: for any subset of families $G \subseteq F$ and any family $f \in G$, $f \in C_\ell^K(G)$ if and only if ℓ can weakly accommodate f alongside $\widehat{G}_\ell^f$.

Claim D.1. C_ℓ^K is envy-free, $C_\ell \subseteq C_\ell^K$, and $C_\ell^T \subseteq C_\ell^K$.

Proof of Claim D.1. We first show that C_ℓ^K is envy-free. Towards a contradiction, suppose that there exists a subset of families $G \subseteq F$ and two families $f \in C_\ell^K(G)$ and $g \in G \setminus C_\ell^K(G)$ such that $g \rhd_\ell f$. Locality ℓ can (weakly) accommodate f alongside $\widehat{G}_\ell^f$ but cannot (weakly) accommodate g alongside $\widehat{G}_\ell^g$. It follows that $\nu_d^f + \sum_{h\in\widehat{G}_\ell^f} \nu_d^h \leq \kappa_\ell^d < \nu_d^g + \sum_{h\in\widehat{G}_\ell^g} \nu_d^h$, a contradiction since $g \rhd_\ell f$ implies that $(\{g\} \cup \widehat{G}_\ell^g) \subseteq \widehat{G}_\ell^f$.

We next show that $C_\ell \subseteq C_\ell^K$. Towards a contradiction, suppose that there exists a family $f \in C_\ell(G) \setminus C_\ell^K(G)$. As $f \notin C_\ell^K(G)$, ℓ cannot (weakly) accommodate f alongside $\widehat{G}_\ell^f$. Then, $f \in C_\ell(G)$ implies that $C_\ell(G)$ is not interference-free, a contradiction.

The proof that $C_\ell^T \subseteq C_\ell^K$ is analogous, one simply needs to replace $C_\ell(G)$ by $C_\ell^T(G)$ in the preceding paragraph. □

Letting $G \subseteq F$ be any subset of families, we show that $|C_\ell(G)| \leq |C_\ell^T(G)|$. A direct implication of Claim D.1 is that the proof is complete if $C_\ell^T(G) = C_\ell^K(G)$ so we focus on the case in which $C_\ell^T(G) \subset C_\ell^K(G)$. Let f be the highest-priority family in $G \setminus C_\ell^T(G)$. As C_ℓ^T is envy-free (Proposition D.1), we have that $C_\ell^T(G) = \widehat{G}_\ell^f$. Moreover, as $C_\ell^T(G) \subset C_\ell^K(G)$ by assumption and $C_\ell^K(G)$ is envy-free by Claim D.1, $f \in C_\ell^K(G)$. It follows that ℓ can (weakly) accommodate f alongside $\widehat{G}_\ell^f$; hence, $\widetilde{\theta}_\ell^f(G) \geq |\widehat{G}_\ell^f(G)| + 1$ (see Algorithm 5). However, as $f \notin C_\ell^T(G)$, we know that $\theta_\ell^f < |\widehat{G}_\ell| + 1$; therefore, there exists a family $g \in \widehat{F}_\ell^f$ such that $\widetilde{\theta}_\ell^g(G) = \theta_\ell^f(G)$. We consider two cases separately.

Case 1: $\widetilde{\theta}_\ell^g(G) = \theta_\ell^f(G) = 0$. By definition (see Algorithm 5), ℓ cannot (weakly) accommodate g alongside $\widehat{G}_\ell^g$. It follows that $g \notin C_\ell^K(G \cup \{g\})$, which implies that $C_\ell^K(G \cup \{g\}) \subseteq \widehat{G}_\ell^g$ since C_ℓ^K is envy-free. Then, by Claim D.1, we have that $C_\ell(G \cup \{g\}) \subseteq \widehat{G}_\ell^g$ so $|C_\ell(G \cup \{g\})| \leq |\widehat{G}_\ell^g|$. As C_ℓ satisfies the cardinal monotonicity condition, it follows that $|C_\ell(G)| \leq |\widehat{G}_\ell^g|$. As $C_\ell^T(G) = \widehat{G}_\ell^f$ and $g \rhd_\ell f$, we have that $\widehat{G}_\ell^g \subseteq C_\ell^T(G)$; hence, we conclude that $|C_\ell(G)| \leq |\widehat{G}_\ell^g| \leq |C_\ell^T(G)|$.

Case 2: $\widetilde{\theta}_\ell^g(G) = \theta_\ell^f(G) > 0$. By definition (see Algorithm 5), there exists a subset of families G^* such that $\widehat{G}_\ell^g \subseteq G^* \subseteq \widehat{F}_\ell^f$ and $|G^*| = \theta_\ell^f(G)$ alongside which ℓ cannot (weakly) accommodate g. Let us define the set $H = G^* \cup G \cup \{g\}$ and observe that, by construction, $G^* = \widehat{H}_\ell^g$. As ℓ cannot (weakly) accommodate g alongside $G^* = \widehat{H}_\ell^g$, we have that $g \notin C_\ell^K(H)$, which implies that $C_\ell^K(H) \subseteq \widehat{H}_\ell^g = G^*$ since C_ℓ^K is envy-free (Claim D.1). By Claim D.1, $C_\ell(H) \subseteq G^*$ so $|C_\ell(H)| \leq |G^*| = \theta_\ell^f(G)$, which implies that $|C_\ell(G)| \leq \theta_\ell^f(G)$ since C_ℓ satisfies the cardinal monotonicity condition. As $\theta_\ell^f(G) < |\widehat{G}_\ell^f| + 1$ (equivalently, $\theta_\ell^f(G) \leq |\widehat{G}_\ell^f|$) and $C_\ell^T(G) = \widehat{G}_\ell^f$, we conclude that $|C_\ell(G)| \leq \theta_\ell^f(G) \leq |C_\ell^T(G)|$.

Having established that $|C_\ell(G)| \leq |C_\ell^T(G)|$, we now show that every family in $C_\ell^T(G) \setminus C_\ell(G)$ has a higher priority than any family in $C_\ell(G) \setminus C_\ell^T(G)$. Considering any two families $f \in C_\ell^T(G) \setminus C_\ell(G)$ and $g \in C_\ell(G) \setminus C_\ell^T(G)$, it is immediate that $f \rhd_\ell g$ since $f \in C_\ell^T(G)$, $g \in G \setminus C_\ell^T(G)$, and $C_\ell^T(G)$ is envy-free (Proposition D.1). □

Proof of Proposition D.4. We first show that the TKDA mechanism matches at least one family to a locality. Recall that we assume throughout that for each 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, $\mu^{\mathrm{TKDA}}(f) = \ell \succ_f \emptyset$; otherwise, f does not propose to ℓ, so $\mu^{\mathrm{TKDA}}(f) \succ_f \ell \succ_f \emptyset$. 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 $|F \setminus \mu^{\mathrm{TKDA}}(\emptyset)| < \min\{|F|, |L| - 1\}$. Then, $\mu^{\mathrm{TKDA}}(\emptyset) \neq \varnothing$ and there exists $\ell \in L \setminus \{\emptyset\}$ such that $\mu^{\mathrm{TKDA}}(\ell) = \varnothing$. Since we assume that $\emptyset$ is every family's last preference and $\mu^{\mathrm{TKDA}}(\ell) = \varnothing$, 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 $f \in \mu^{\mathrm{TKDA}}(\ell)$, which contradicts $\mu^{\mathrm{TKDA}}(\ell) = \varnothing$.

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

Proof of Proposition D.5. 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 matching constraints (see the proof of Theorem 6), the same is true of the TKDAC algorithm, i.e., for every $\ell \in L$, ℓ can accommodate $\mu^{\mathrm{TKDAC}}(\ell)$.

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 $\widetilde{\triangleright}^0 = \triangleright$. For each Step $j = 0, 1, \ldots, N$, each family $f \in F$, and each locality $\ell \in L$, let

- $\widehat{F}_\ell^f(\widetilde{\triangleright}^j)$ be the set of families that are higher than f on $\widetilde{\triangleright}_\ell^j$;
- 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 $\widehat{F}_\ell^f(\widetilde{\triangleright}^0) = \widehat{F}_\ell^f(\widetilde{\triangleright}^1) = \widehat{F}_\ell^f(\triangleright) = \widehat{F}_\ell^f$ and $\widehat{F}_\ell^f(\widetilde{\triangleright}^N) = \widehat{F}_\ell^f(\widetilde{\triangleright})$. Since no rejections, proposals, or clinches occur in Step 0, we also have $\Delta_f^0 = \Gamma_f^0 = \Theta_\ell^0 = \varnothing$.

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

Lemma D.1. For each Step $j = 1, \ldots, N$, each family $f \in F$, and each locality $\ell \in L$:

(i) If $f \notin \Theta_{\ell'}^{j-1}$ for all $\ell' \in L$ such that $\ell' \succ_f \ell$, then $\widehat{F}_\ell^f(\widetilde{\triangleright}^j) \subseteq \widehat{F}_\ell^f(\widetilde{\triangleright}^{j-1})$;

(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) $\Theta_\ell^{j-1} \subseteq \Theta_\ell^j$.

(The proof of Lemma D.3 follows at the end of the proof of Proposition D.5.) Part (i) of Lemma D.1 says that if family f has not clinched a locality that f prefers to ℓ, then the set of families that 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 cliched ℓ grows throughout the Clinching Round.

Lemma D.2. For each Step $j = 1, \ldots, N$, each family $f \in F$, and each locality $\ell \in L$:

(i) If $\ell \in \Delta_f^j$, then $\mu^{\text{TKDAC}}(f) \neq \ell$;

(ii) If $\ell \in \Gamma_f^j$, then $\mu^{\text{TKDAC}}(f) \succeq_f \ell$; and

(iii) If $f \in \Theta_\ell^j$, then $\mu^{\text{TKDAC}}(f) = \ell$.

(The proof of Lemma D.3 follows at the end of the proof of Proposition D.5.) Part (i) of Lemma D.2 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 D.1 and D.2 to show that the TKDAC mechanism is strategy-proof and interference-free and that $\mu^{\text{TKDAC}} \succeq \mu^{\text{TKDA}}$.

Proof that 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, $\widetilde{\triangleright}^j$, 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 $\overline{N}$, $\widetilde{\triangleright}^j$, $\overline{\Delta}_f^j$, $\overline{\Gamma}_f^j$, and $\overline{\Theta}_\ell^j$ 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 $j = 1, \ldots, m-1$ and $f \in \cup_{\ell \in L}\{\Theta_\ell^j\}$ for all $j = m, \ldots, N$. (By Lemma D.1(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 $\overline{m}$ 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 D.3. Let $q = \min\{m, \overline{m}, N, \overline{N}\}$. For each $j = 1, \ldots, q$, each $g \in F$, and each $\ell \in L$,

$$\Theta_\ell^j \setminus \{f\} = \overline{\Theta}_\ell^j \setminus \{f\}, \quad \Delta_g^j = \overline{\Delta}_g^j, \quad \Gamma_g^j = \overline{\Gamma}_g^j, \quad \text{and} \quad \widetilde{\rhd}^j = \overline{\rhd}^j.$$

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

(The proof of Lemma D.3 follows at the end of the proof of Proposition D.5.) There are four cases to consider.

Case 1: $m = \infty$ and $\overline{m} = \infty$. In this case, f does not clinch any locality in the Clinching Round, irrespective of whether f reports $\succ_f$ or $\succ'_f$. Then, $q = \min\{N, \overline{N}\} < \min\{m, \overline{m}\}$. If $N \leq \overline{N}$, 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 D.3, $\Theta_\ell^{N-1} = \overline{\Theta}_\ell^{N-1}$ and $\Theta_\ell^N = \overline{\Theta}_\ell^N$ for all $\ell \in L$; therefore, $\overline{\Theta}_\ell^{N-1} = \overline{\Theta}_\ell^N$ for all $\ell \in L$, which means that $N = \overline{N}$. Then, since $\widetilde{\rhd}^N = \overline{\rhd}^N$ by Lemma D.3, we conclude that $\widetilde{\rhd} = \overline{\rhd}$. 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 $\widetilde{\rhd}$ 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 $\overline{\rhd} = \widetilde{\rhd}$ respectively. Since the TKDA mechanism 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 $N \geq \overline{N}$.

Case 2: $m \leq \min\{\overline{m}, N\}$. 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 $\overline{m} \geq m$ or does not clinch any locality. Since $m \leq \min\{\overline{m}, N\}$, $q = \min\{m, \overline{N}\}$. Towards a contradiction, suppose that $m > \overline{N}$. Then, $q = \overline{N} < \min\{m, \overline{m}\}$. As the Clinching Round ends whenever the same clinches occur in two consecutive rounds, we have that $\Theta_\ell^{\overline{N}-1} = \Theta_\ell^{\overline{N}}$ for all $\ell \in L$. Moreover, Lemma D.3 implies that $\Theta_\ell^{\overline{N}-1} = \overline{\Theta}_\ell^{\overline{N}-1}$ and $\Theta_\ell^{\overline{N}} = \overline{\Theta}_\ell^{\overline{N}}$ for all $\ell \in L$. It follows that $\overline{\Theta}_\ell^{\overline{N}-1} = \overline{\Theta}_\ell^{\overline{N}}$, so $N = \overline{N}$, which contradicts our assumption that $\overline{N} < m \leq N$.

We have established that $m \leq \overline{N}$; 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 D.2(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 D.3, $\Delta_f^m = \overline{\Delta}_f^m$; therefore

$\ell' \in \overline{\Delta}_f^m$ for all $\ell' \in L$ such that $\ell' \succ_f \ell$. By Lemma D.2(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: $\overline{m} \leq \min\{m, \overline{N}\}$. In this case, if f reports $\succ'_f$, then f clinches a locality in Step $\overline{m}$ of the Clinching Round and, if f reports truthfully, then f either clinches a locality in Step $m \geq \overline{m}$ or does not clinch any locality. Since $\overline{m} \leq \{m, \overline{N}\}$, $q = \min\{\overline{m}, N\}$. Using analogous reasoning to Case 2, we establish that $\overline{m} \leq N$; hence $q = \overline{m}$. When f reports $\succ'_f$, f clinches a locality denoted by ℓ in Step $\overline{m}$ of the Clinching Round, i.e., $f \in \overline{\Theta}_\ell^{\overline{m}}$. By Lemma D.2(iii), we have that $\varphi^{TKDAC}(\succ'_f, \succ_{-f}) = \ell$. By construction (Step $\overline{m}$(c) of the Clinching Round), $f \in \overline{\Theta}_\ell^{\overline{m}}$ implies that $\ell \in \overline{\Gamma}_f^{\overline{m}}$. By Lemma D.3, $\Gamma_f^{\overline{m}} = \overline{\Gamma}_f^{\overline{m}}$ so $\ell \in \Gamma_f^{\overline{m}}$; therefore, by Lemma D.2(ii), $\varphi^{TKDAC}(\succ)(f) \succeq_f \ell$, as required.

Proof that TKDAC is interference-free. Arbitrarily fix a family $f \in F$ and, for ease of notation, let $\ell = \mu^{\text{TKDAC}}(f)$. We need to show that f does not interfere with μ^{TKDAC}. Towards a contradiction, suppose that f interferes with μ^{TKDAC}; that is, ℓ cannot weakly accommodate f alongside $\widehat{F}^f_{\mu^{\text{TKDAC}}} = \{g \in F : g \rhd_\ell f \text{ and } \ell \succeq_g \mu^{\text{TKDAC}}(g)\}$. As μ^{TKDAC} is the outcome of the TKDA mechanism with the priority profile $\widetilde{\rhd}$ and the TKDA mechanism is interference-free, ℓ can weakly accommodate f alongside $\{g \in F : g \,\widetilde{\rhd}_\ell f \text{ and } \ell \succeq_g \mu^{\text{TKDAC}}(g)\}$. Therefore, there exists a family

$$h \in \{g \in F : g \rhd_\ell f \text{ and } \ell \succeq_g \mu^{\text{TKDAC}}(g)\} \setminus \{g \in F : g \,\widetilde{\rhd}_\ell f \text{ and } \ell \succeq_g \mu^{\text{TKDAC}}(g)\}.$$

By assumption, $h \rhd_\ell f$ and $f \,\widetilde{\rhd}_\ell\, h$; equivalently, $f \in \widehat{F}^h_\ell(\widetilde{\rhd}) \setminus \widehat{F}^h_\ell$. Therefore, there exists a Step $j = 1, \ldots, N$ of the Clinching Round such that $f \in \widehat{F}^h_\ell(\widetilde{\rhd}^j) \setminus \widehat{F}^h_\ell(\widetilde{\rhd}^{j-1})$. By the contrapositive of Lemma D.1(i), it follows that $h \in \Theta^{j-1}_{\ell'}$ for some $\ell' \in L$ such that $\ell' \succ_h \ell$. By Lemma D.2(iii), we have that $\mu^{\text{TKDAC}}(h) = \ell'$. We conclude that $\mu^{\text{TKDAC}}(h) \succ_h \ell$, a contradiction since, by definition, $\ell \succeq_h \mu^{\text{TKDAC}}(h)$.

Proof that TKDAC is weakly more efficient than TKDA ($\mu^{\text{TKDAC}} \succeq \mu^{\text{TKDA}}$).

We first introduce some additional notation. Let M be the number of rounds of the TKDA algorithm and, for each Round $i = 1, \ldots, M$ and each locality $\ell \in L$, let Π^i_ℓ 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 $\widetilde{\rhd}$ and lasts $\widetilde{M}$ rounds. For each Round $i = 1, \ldots, \widetilde{M}$ and each locality $\ell \in L$, let $\widetilde{\Pi}^i_\ell$ 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, $\theta^f_\ell(\Pi^i_\ell \cap \widehat{F}^f_\ell)$ is the threshold of family f for locality ℓ in Round i of the TKDA algorithm and $\theta^f_\ell(\widetilde{\Pi}^i_\ell \cap \widehat{F}^f_\ell(\widetilde{\rhd}))$ is the threshold of family f for locality ℓ in Round i of the TKDAC algorithm.

Claim D.2. In every Round $i = 1, \ldots, \widetilde{M}$ of the TKDAC algorithm, every family $f \in F$ proposes to a locality that f weakly prefers to $\mu^{\text{TKDA}}(f)$.

By construction, for every family $f \in F$, $\mu^{\text{TKDAC}}(f)$ is the locality to which f proposes in Round $\widetilde{M}$ of the TKDAC algorithm. Therefore, Claim D.2 implies that $\mu^{\text{TKDAC}}(f) \succeq_f \mu^{\text{TKDA}}(f)$ for all $f \in F$, as required. To complete the proof of Proposition D.5 it only remains to prove Claim D.2 as well as Lemmata D.1, D.2 and D.3. ☐

Proof of Claim D.2. In Round 1 of the TKDAC algorithm, every family proposes to its most preferred locality; therefore Claim D.2 holds for $i = 1$. The remainder of the proof is by induction. We suppose that Claim D.2 holds for some $i = 1, \ldots, \widetilde{M}-1$ (induction hypothesis) and show that Claim D.2 holds for $i+1$.

Consider any family $f \in F$ and, for ease of notation, let $\ell = \mu^{\text{TKDA}}(f)$. 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

$$\theta_\ell^f(\widetilde{\Pi}_\ell^i \cap \widehat{F}_\ell^f) \geq |\widetilde{\Pi}_\ell^i \cap \widehat{F}_\ell^f| + 1. \tag{24}$$

First, as f proposes to ℓ in Round i, $\ell \succeq_f \mu^{\text{TKDAC}}(f)$ by construction. Then, by Lemma D.2(iii), f does not clinch any locality that f strictly prefers to ℓ in the Clinching Round. By Lemma D.1(i), it follows that $\widehat{F}_\ell^f(\widetilde{\rhd}) \subseteq \widehat{F}_\ell^f$.

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

$$\theta_\ell^f(\Pi_\ell^M \cap \widehat{F}_\ell^f) \geq |\Pi_\ell^M \cap \widehat{F}_\ell^f| + 1. \tag{25}$$

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

Case 1: $\theta_\ell^f(\Pi^M \cap \widehat{F}_\ell^f) = \infty$. In this case, by the definition of thresholds (Algorithm 5), ℓ can weakly accommodate f alongside $\widehat{F}_\ell^f$. As $\widehat{F}_\ell^f(\widetilde{\rhd}) \subseteq \widehat{F}_\ell^f$, ℓ can weakly accommodate f alongside $\widehat{F}_\ell^f(\widetilde{\rhd})$. Again, by the definition of thresholds (Algorithm 5), we have that $\theta_\ell^f(\widetilde{\Pi}^i \cap \widehat{F}_\ell^f(\widetilde{\rhd})) = \infty$; hence inequality (24) is satisfied, as required.

Case 2: $\theta_\ell^f(\Pi^M \cap \widehat{F}_\ell^f) \neq \infty$. 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: Locality 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 $j = 1, \ldots, M$ of the TKDA algorithm, $g \in \Pi_\ell^j$ and

$$\theta_\ell^g(\Pi_\ell^j \cap \widehat{F}_\ell^g) < |\Pi_\ell^j \cap \widehat{F}_\ell^g| + 1. \tag{26}$$

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:

$$\theta_\ell^g(\Pi_\ell^M \cap \widehat{F}_\ell^g) < |\Pi_\ell^M \cap \widehat{F}_\ell^g| + 1. \tag{27}$$

First, inequality (27) holds trivially if $\theta_\ell^g(\Pi_\ell^M \cap \widehat{F}_\ell^g) = 0$. Second, inequality (26) implies that $\theta_\ell^g(\Pi_\ell^j \cap \widehat{F}_\ell^g) \neq \infty$; hence Lemma A.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 A.1(ii) to obtain that

$$\theta_\ell^g(\Pi_\ell^M \cap \widehat{F}_\ell^g) \leq \theta_\ell^g(\Pi_\ell^j \cap \widehat{F}_\ell^g) + |\Pi_\ell^M \cap \widehat{F}_\ell^g| - |\Pi_\ell^j \cap \widehat{F}_\ell^g|. \tag{28}$$

Combined with inequality (26), inequality (28) implies inequality (27).

As $g \rhd_\ell f$, 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 (27) and (25) yields

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

By inequality (26), $\theta_\ell^g(\Pi_\ell^M \cap \widehat{F}_\ell^f) \neq \infty$; hence $\theta_\ell^g(\Pi_\ell^M \cap \widehat{F}_\ell^g) \neq \infty$. Hence, we can apply Lemma A.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 (29). We therefore conclude that Sub-case 2.1 cannot occur.

Sub-case 2.2: Locality 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 $\mu^{\text{TKDA}}(g)$; therefore, the fact that $g \in \widetilde{\Pi}_\ell^i$ implies that $\ell \succeq_g \mu^{\text{TKDA}}(g)$. 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 $\ell \neq \mu^{\text{TKDA}}(g)$. 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 $\mu^{\text{TKDA}}(g) \succeq_g \ell$. Combining these observations, we have that $\ell \succeq_g \mu^{\text{TKDA}}(g)$, $\ell \neq \mu^{\text{TKDA}}(g)$, and $\mu^{\text{TKDA}}(g) \succeq_g \ell$, a contradiction.

Having established that $\widetilde{\Pi}^i_\ell \cap \widehat{F}^f_\ell \subseteq \Pi^M_\ell \cap \widehat{F}^f_\ell$ and that $\widehat{F}^f_\ell(\widetilde{\triangleright}) \subseteq \widehat{F}^f_\ell$, we conclude that $\widetilde{\Pi}^i_\ell \cap \widehat{F}^f_\ell(\widetilde{\triangleright}) \subseteq \Pi^M_\ell \cap \widehat{F}^f_\ell$. It follows that $|\widetilde{\Pi}^i_\ell \cap \widehat{F}^f_\ell(\widetilde{\triangleright})| \leq |\Pi^M_\ell \cap \widehat{F}^f_\ell|$. Moreover, as $\theta^f_\ell(\Pi^M_\ell \cap \widehat{F}^f_\ell) \neq \infty$ by the assumption of Case 2, and as $\theta^f_\ell(\Pi^M_\ell \cap \widehat{F}^f_\ell) \neq 0$ by inequality (25), we have that $\theta^f_\ell(\Pi^M_\ell \cap \widehat{F}^f_\ell) \in \mathbb{Z}_{>0}$ so we can apply Lemma A.1(ii) to obtain that

$$\theta^f_\ell(\Pi^M_\ell \cap \widehat{F}^f_\ell) \leq \theta^f_\ell(\widetilde{\Pi}^i_\ell \cap \widehat{F}^f_\ell(\widetilde{\triangleright})) + |\Pi^M_\ell \cap \widehat{F}^f_\ell| - |\widetilde{\Pi}^i_\ell \cap \widehat{F}^f_\ell(\widetilde{\triangleright})|. \tag{30}$$

Combining inequality (30) with inequality (25) yields inequality (24), as required. □

Proof of Lemma D.1. We prove the lemma by a single induction argument. To show that the lemma holds for $j = 1$, note that (i) $\widetilde{\triangleright}^0 = \widetilde{\triangleright}^1 = \triangleright$, (ii) $\varnothing = \Delta^0_f \subseteq \Delta^1_f$ for every $f \in F$, (iii) $\varnothing = \Gamma^0_f \subseteq \Gamma^1_f$ for every $f \in F$, and (iv) $\varnothing = \Theta^0_\ell \subseteq \Theta^1_\ell$ for every $\ell \in L$.

For the induction step, let us assume that part (iv) of the lemma holds for some $j = 1, \ldots, N$, i.e., $\Theta^{j-1}_\ell \subseteq \Theta^j_\ell$ for every $\ell \in L$ and every $j = 1, \ldots, N$. 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 $f \notin \Theta^j_{\ell'}$ for all $\ell' \in L$ such that $\ell' \succ_f \ell$, then $\widehat{F}^f_\ell(\widetilde{\triangleright}^{j+1}) \subseteq \widehat{F}^f_\ell(\widetilde{\triangleright}^j)$;
(ii) $\Delta^j_f \subseteq \Delta^{j+1}_f$;
(iii) If $\ell \in \Gamma^j_f \setminus \Gamma^{j+1}_f$, then there exists $\ell' \in L$ such that $\ell' \succ_f \ell$ and $f \in \Theta^j_{\ell'}$; and
(iv) $\Theta^j_\ell \subseteq \Theta^{j+1}_\ell$.

Proof of (i). We prove the contrapositive. Suppose there exists a family $g \in \widehat{F}^f_\ell(\widetilde{\triangleright}^{j+1}) \setminus \widehat{F}^f_\ell(\widetilde{\triangleright}^j)$, then $f \,\widetilde{\triangleright}^j_\ell\, g$ and $g \,\widetilde{\triangleright}^{j+1}_\ell\, f$. 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 $g \,\widetilde{\triangleright}^{j+1}_\ell\, f$, 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 $f \,\widetilde{\triangleright}^j_\ell\, g$, 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 $\Theta^{j-1}_\ell \subseteq \Theta^j_\ell$, g continues to clinch that locality in Step j(c) of the Clinching Round. Since $g \,\widetilde{\triangleright}^{j+1}_\ell\, f$, 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 $\ell \in \Delta^j_f$, we need to show that $\ell \in \Delta^{j+1}_f$. As $\ell \in \Delta^j_f$, ℓ cannot weakly accommodate f alongside $\Theta^{j-1}_\ell \cap \widehat{F}^f_\ell(\widetilde{\triangleright}^j)$. Similarly, $\ell \in \Delta^{j+1}_f$ if ℓ cannot weakly accommodate f alongside $\Theta^j_\ell \cap \widehat{F}^f_\ell(\widetilde{\triangleright}^{j+1})$. Therefore, it is sufficient to show that

$$(\Theta^{j-1}_\ell \cap \widehat{F}^f_\ell(\widetilde{\triangleright}^j)) \subseteq (\Theta^j_\ell \cap \widehat{F}^f_\ell(\widetilde{\triangleright}^{j+1})).$$

Towards a contradiction, suppose that there exists a family

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

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

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

Since $g \in \widehat{F}_\ell^f(\widetilde{\triangleright}_\ell^j) \setminus \widehat{F}_\ell^f(\widetilde{\triangleright}^{j+1})$, we have that $g \,\widetilde{\triangleright}_\ell^j f$ and $f \,\widetilde{\triangleright}_\ell^{j+1} g$. Therefore, $f \in \widehat{F}_\ell^g(\widetilde{\triangleright}_\ell^{j+1}) \setminus \widehat{F}_\ell^g(\widetilde{\triangleright}^j)$ and, as a result, $\widehat{F}_\ell^g(\widetilde{\triangleright}_\ell^{j+1}) \not\subseteq \widehat{F}_\ell^g(\widetilde{\triangleright}^j)$. 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 $\widehat{F}_\ell^f(\widetilde{\triangleright}^j)$ but not alongside $\widehat{F}_\ell^f(\widetilde{\triangleright}^{j+1})$, which implies that $\widehat{F}_\ell^f(\widetilde{\triangleright}^{j+1}) \not\subseteq \widehat{F}_\ell^f(\widetilde{\triangleright}^j)$. 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., $f \in \Theta_\ell^{j+1}$. □

Proof of Lemma D.2. For ease of notation, let $\mu = \mu^{\text{TKDAC}}$. We consider a Step $j = 1, \ldots, N$ 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 $i = 2, \ldots, N$ 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 $i = 2, \ldots, N$ of the Clinching Round. Since $\mu(g)$ rejects g in Step i(a) of the Clinching Round, $\mu(g)$ cannot weakly accommodate g alongside $\Theta_{\mu(g)}^{i-1} \cap \widehat{F}_{\mu(g)}^g(\widetilde{\triangleright}^i)$. If all families in $\Theta_{\mu(g)}^{i-1} \cap \widehat{F}_{\mu(g)}^g(\widetilde{\triangleright}^i)$ are matched to $\mu(g)$ at the end of the TKDAC algorithm (i.e., if $(\Theta_{\mu(g)}^{i-1} \cap \widehat{F}_{\mu(g)}^g(\widetilde{\triangleright}^i)) \subseteq \mu(\mu(g))$), then $\mu(g)$ can accommodate $(\Theta_{\mu(g)}^{i-1} \cap \widehat{F}_{\mu(g)}^g(\widetilde{\triangleright}^i)) \cup \{g\}$. Therefore, $\mu(g)$ can weakly accommodate g alongside $\Theta_{\mu(g)}^{i-1} \cap \widehat{F}_{\mu(g)}^g(\widetilde{\triangleright}^i)$, a contradiction. We conclude that there exists a family $h \in \Theta_{\mu(g)}^{i-1} \cap \widehat{F}_{\mu(g)}^g(\widetilde{\triangleright}^i)$ 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 D.1(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 $\widehat{F}^h_\ell(\widetilde{\triangleright}^N) = \widehat{F}^h_\ell(\widetilde{\triangleright})$. 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 $h \in \Theta^{i-1}_{\mu(g)}$ 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 $\mu(f) \in \Delta^j_f$. The preceding argument implies, by induction, that there exists a family f' such that $\mu(f') \in \Delta^1_{f'}$, i.e., $\mu(f')$ rejects f' in Step 1 of the Clinching Round. Then, $\mu(f')$ cannot weakly accommodate f' alongside $\Theta^0_{\mu(f')} \cap \widehat{F}^{f'}_{\mu(f')}(\widetilde{\triangleright}^1)$. Since $\Theta^0_{\mu(f')} = \varnothing$, 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 $\ell \in \Gamma^j_f$. We need to show that $\mu(f) \succeq_f \ell$. Then, by Lemma D.1(iii), $\ell \in \Gamma^j_f$ implies that one of the following two cases holds: either $\ell \in \Gamma^N_f$ or $f \in \Theta^{j-1}_{\ell'}$ for some $\ell' \in L$ such that $\ell' \succ_f \ell$.

Case 1: $\ell \in \Gamma^N_f$. In this case, by construction (Step N(c) of the Clinching Round), ℓ can weakly accommodate f alongside $\widehat{F}^f_\ell(\widetilde{\triangleright}^N) = \widehat{F}^f_\ell(\widetilde{\triangleright})$. 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: $f \in \Theta^{j-1}_{\ell'}$ for some $\ell' \in L$ such that $\ell' \succ_f \ell$. In this case, Lemma D.1(iv), we have that $f \in \Theta^N_{\ell'}$ 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., $\ell' \in \Gamma^N_f$. 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 $f \in \Theta^j_\ell$. Then, by construction (Step j(c) of the Clinching Round) we have that $\ell \in \Gamma^j_f$ and, for all $\ell' \in L$ such that $\ell' \succ_f \ell$, it is the case that $\ell' \in \Delta^j_f$. 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 D.3. We prove the first part of the lemma by a single induction argument. For the initial step, we have that $\Theta^0_\ell = \overline{\Theta}^0_\ell = \varnothing$ for all $\ell \in L$.

For the induction step, let us assume that $\Theta^{j-1}_\ell \setminus \{f\} = \overline{\Theta}^{j-1}_\ell \setminus \{f\}$ holds for some $j = 1, \ldots, q$ and all $\ell \in L$. We will show that $\widetilde{\triangleright}^j = \overline{\triangleright}^j$, $\widetilde{\Delta}^j_g = \overline{\Delta}^j_g$, $\widetilde{\Gamma}^j_g = \overline{\Gamma}^j_g$, and, finally, that $\Theta^j_\ell \setminus \{f\} = \overline{\Theta}^j_\ell \setminus \{f\}$ 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, $\Theta^{j-1}_\ell \setminus \{f\} = \overline{\Theta}^{j-1}_\ell \setminus \{f\}$ for all $\ell \in L$ implies that $\Theta^{j-1}_\ell = \overline{\Theta}^{j-1}_\ell$ for all $\ell \in L$. In Step $j-1$(e) of the Clinching Round the construction of $\widetilde{\rhd}^j$ 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 $\widetilde{\rhd}^j = \overline{\rhd}^j$. In Step j(a), the fact that $\Theta_\ell^{j-1} = \overline{\Theta}_\ell^{j-1}$ for all $\ell \in L$ and $\widetilde{\rhd}^j = \overline{\rhd}^j$ 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), $\widetilde{\rhd}^j = \overline{\rhd}^j$ 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 $j = 1, \ldots, q$ 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 $\Theta_\ell^j = \overline{\Theta}_\ell^j$ for all $\ell \in L$, as required. $\square$

Appendix E. Additional Examples

We present an illustrative example of the KTTCE algorithm in Online Appendix E.1 and of the TKDA algorithm in Online Appendix E.2. In Online Appendix E.3, we revisit the latter example to show how clinching can improve the efficiency of the TKDA algorithm.

E.1. 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:

$$\rhd_{\ell_1}: f_1, f_2, \ldots \quad \rhd_{\ell_2}: \boldsymbol{f_1}, \boldsymbol{f_2}, \ldots \quad \rhd_{\ell_3}: \boldsymbol{f_3}, \ldots \quad \rhd_{\ell_4}: \boldsymbol{f_4}, \ldots.$$

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 E.1. In Round 1, the unique trading cycle is $f_1 \to \ell_3 \to f_3 \to \ell_2 \to f_1$. That 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 $f_3 \to \ell_3 \to f_3$ 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, $\mu^4(\ell_2) = \varnothing$ since f_1 has left ℓ_2 to be permanently matched to ℓ_1. As a result, the cycle $f_2 \to \ell_4 \to f_4 \to \ell_2 \to f_2$ 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 E.2. As for KTTCE3, the unique cycle in Round 1 is $f_1 \to \ell_3 \to f_3 \to \ell_2 \to f_1$, 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 $f_1 \to \ell_3 \to f_3 \to \ell_2 \to f_1$, 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 point 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 $\ell_3, \boldsymbol{\ell_2}, \ldots$, KTTCE3 produces

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

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 $f_2 \to \ell_4 \to f_4 \to \ell_2 \to f_2$ 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 $f_1 \to \ell_3 \to f_3 \to \ell_2 \to f_1$ 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 $\rhd_{\ell_1} : f_1, f_2, \ldots$ and $\rhd_{\ell_2} : f_2, f_1, \ldots$ (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 $\rhd_{\ell_1} : f_2, f_1, \ldots$ and $\rhd_{\ell_2} : f_2, f_1, \ldots$ or $\rhd_{\ell_1} : f_2, f_1, \ldots$ and $\rhd_{\ell_2} : f_1, f_2, \ldots$, KTTCE3 and KTTCE4 respectively produce the following matchings:

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

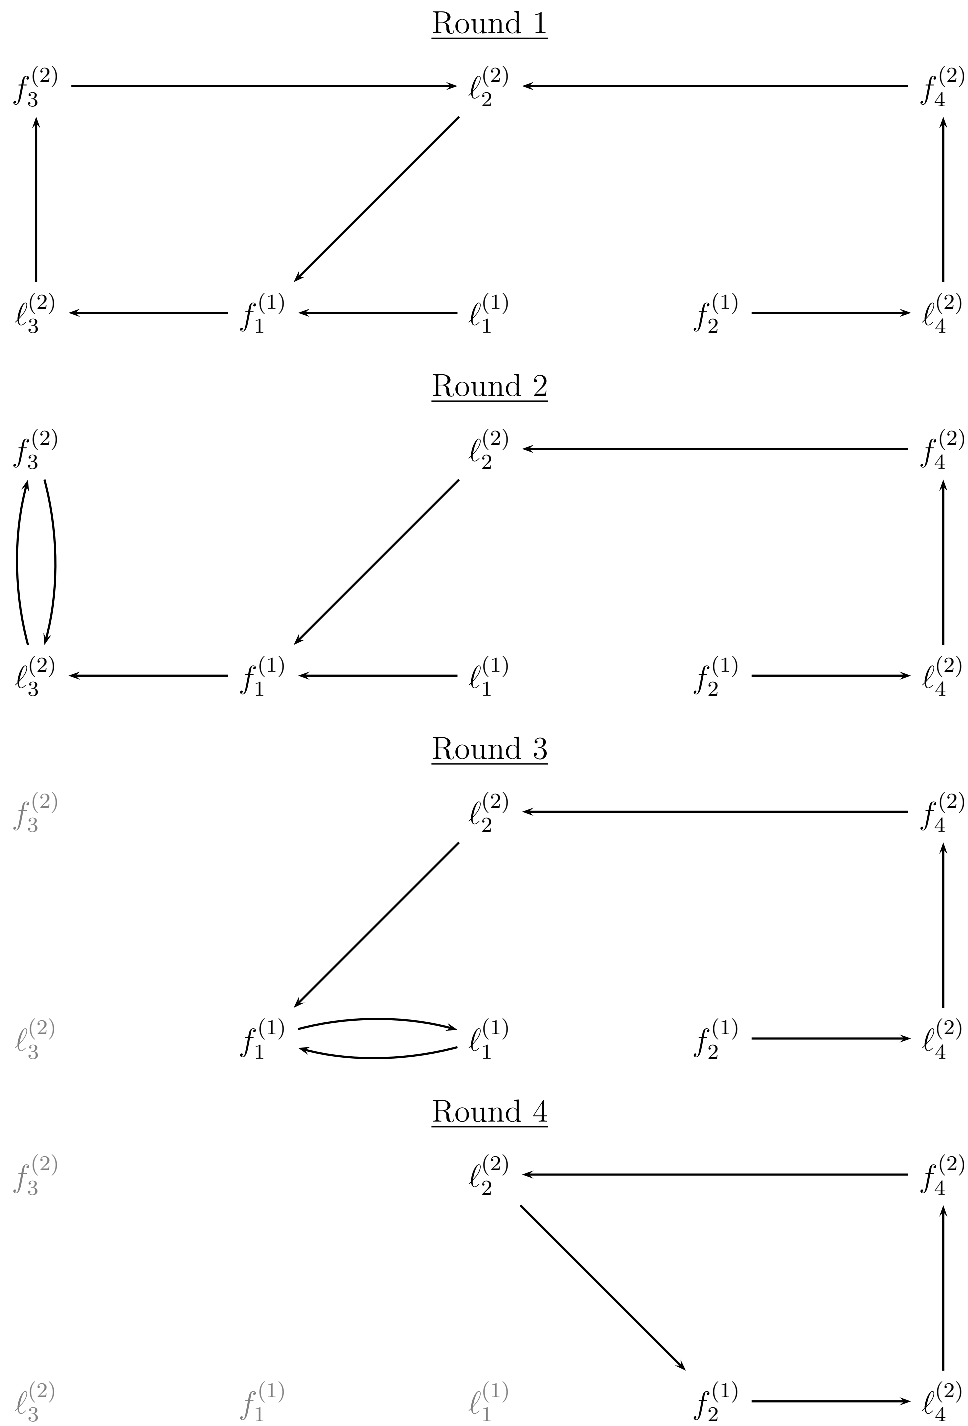

Figure E.1. Workings of KTTC3.

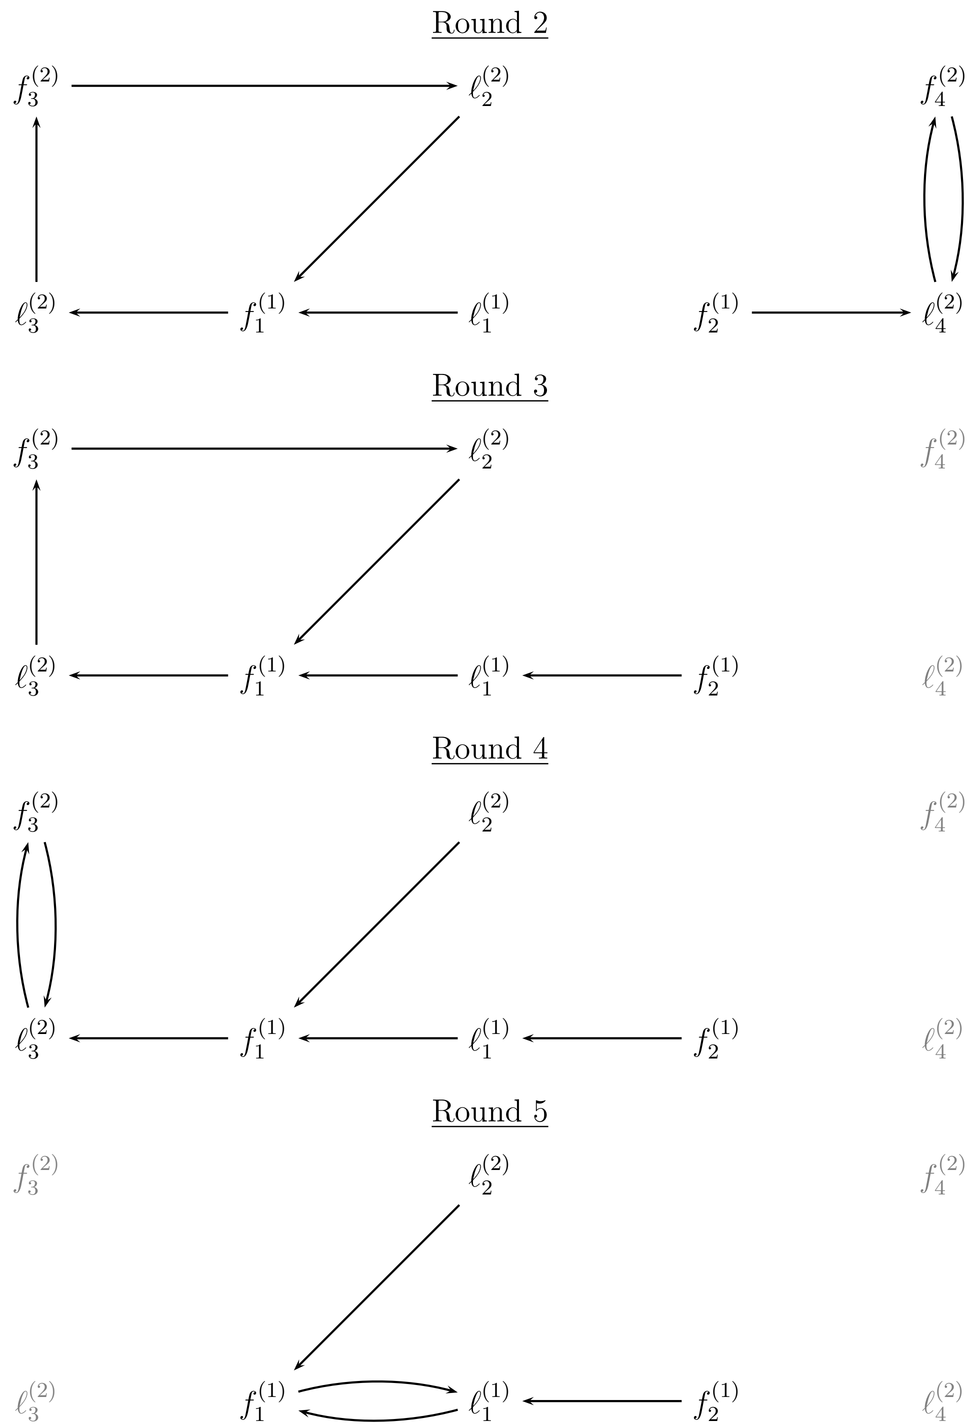

FIGURE E.2. Workings of KTTCE4.

E.2. **Threshold Knapsack Deferred Acceptance.** We illustrate the TKDA algorithm with the following example. There are seven families, four localities, and two dimensions.

Preferences:

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

Priorities:

$$\rhd_{\ell_1} : f_1, f_2, f_3, f_4, f_5, f_6, f_7 \quad \rhd_{\ell_2} : f_5, f_1, f_2, f_7, f_6, f_3, f_4$$
$$\rhd_{\ell_3} : f_4, f_6, f_1, \dots \quad \rhd_{\ell_4} : f_7, f_2, \dots.$$

Sizes and capacities:

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

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

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

The TKDA algorithm lasts four rounds. The first three rounds are displayed in Table E.1. In Round 1, every family proposes to its first-preference locality. Families f_4, f_5, and f_6 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, an interference-free choice function that allowed this would violate the cardinal monotonicity condition. 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, interference-freeness dictates that ℓ_1 must only tentatively accept two families: f_1 and f_2. However, the cardinal monotonicity condition dictates that at most two families can be tentatively accepted when f_4, f_5, and f_6 propose.

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	$(2,0)$	∞	$\boxed{f_7}$	$(1,0)$	∞
f_2	$(1,0)$	∞	$\boxed{f_1}$	$(2,1)$	1	f_6	$(1,0)$	∞	$\boxed{\cancel{f_2}}$	$(1,0)$	0
f_3	$(1,0)$	∞	f_2	$(1,0)$	∞	f_1	$(2,1)$	1	$\vdots$		
$\boxed{f_4}$	$(2,0)$	2	f_7	$(1,0)$	1	$\vdots$					
$\boxed{f_5}$	$(1,1)$	2	f_6	$(1,0)$	1						
$\boxed{\cancel{f_6}}$	$(1,0)$	2	$\boxed{\cancel{f_3}}$	$(1,0)$	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	$(2,0)$	∞	$\boxed{f_7}$	$(1,0)$	∞
$\boxed{f_2}$	$(1,0)$	∞	$\boxed{f_1}$	$(2,1)$	1	f_6	$(1,0)$	∞	f_2	$(1,0)$	0
$\boxed{f_3}$	$(1,0)$	∞	f_2	$(1,0)$	∞	f_1	$(2,1)$	1	$\vdots$		
$\boxed{f_4}$	$(2,0)$	3	f_7	$(1,0)$	1	$\vdots$					
$\boxed{\cancel{f_5}}$	$(1,1)$	0	$\boxed{\cancel{f_6}}$	$(1,0)$	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	$(2,1)$	∞	$\boxed{f_5}$	$(1,1)$	∞	f_4	$(2,0)$	∞	$\boxed{f_7}$	$(1,0)$	∞
$\boxed{f_2}$	$(1,0)$	∞	$\boxed{\cancel{f_1}}$	$(2,1)$	0	$\boxed{f_6}$	$(1,0)$	∞	f_2	$(1,0)$	0
$\boxed{f_3}$	$(1,0)$	∞	f_2	$(1,0)$	∞	f_1	$(2,1)$	2	$\vdots$		
$\boxed{f_4}$	$(2,0)$	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 E.1. Rounds 1-3 of the TKDA algorithm. Sizes and capacities in parentheses. $\boxed{f_i}$: f_i proposes and is tentatively accepted. $\boxed{\cancel{f_i}}$: f_i proposes and is permanently rejected.

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 interference-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 at 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.

E.3. **Threshold Knapsack Deferred Acceptance with Clinching.** We use the example introduced in Online Appendix E.2 to illustrate the TKDAC algorithm. We show that

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

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

The three steps of the Clinching Round are displayed in Table E.2. 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), $\widetilde{\triangleright}^2$ is constructed by giving f_7 the lowest priority at ℓ_1, ℓ_2, ℓ_3. In particular, ℓ_2's priority list is updated to

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

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), $\widetilde{\triangleright}^3$ is constructed by moving f_2 to the bottom of ℓ_2 and ℓ_3's priorities. f_2's priority list is updated to

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

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

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	$(2,0)$	✓	$\boxed{f_7}$	$(1,0)$	✓
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	$(2,0)$	✓	$\boxed{f_7}$	$(1,0)$	✓
$\boxed{f_2}$	$(1,0)$	✓	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	$(2,0)$	✓	$\boxed{f_7}$	$(1,0)$	✓
$\boxed{f_2}$	$(1,0)$	✓	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 E.2. Clinching Round of the TKDAC algorithm. Sizes and capacities in parentheses. ✓: a locality proposes to a family. ✗: a locality rejects a family. $\boxed{f_i}$: family f_i clinches a locality.

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 $\widetilde{\rhd} = \widetilde{\rhd}^3$ such that

$$\rhd_{\ell_1} : f_1, f_2, f_3, f_4, f_5, f_6, f_7 \quad \rhd_{\ell_2} : f_5, f_1, f_6, f_3, f_4, f_2, f_7$$
$$\rhd_{\ell_3} : f_4, f_6, f_1, \ldots \quad \rhd_{\ell_4} : f_7, f_2, \ldots$$

With the priority profile $\widetilde{\rhd}$ constructed in the Clinching Round, the TKDA algorithm lasts four rounds. The first three rounds are displayed in Table E.3. Changing the priority profile from $\rhd$ to $\widetilde{\rhd}$ affects the TKDA algorithm in one important way. With the constructed priority profile $\widetilde{\rhd}$, 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 $\mu^{\text{TKDAC}}(f_6) = \ell_2$.

When the TKDA algorithm is run with the true priority profile $\rhd$, 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 the interference-freeness or cardinal monotonicity conditions when f_6 is matched to ℓ_2.

While the preceding example has two categories, clinching can also yield efficiency gains when $|D| = 1$. Consider a simple example with three families f_1 (size 2), f_2 (size 1), and f_3 (size 1) and two localities ℓ_1 (capacity 2) and ℓ_2 (capacity 2). All families prefer ℓ_1 to ℓ_2 and ℓ_2 to the null. Both localities prioritize f_1 over f_2 and f_2 over f_3.

It is easy to verify that the TKDA algorithm matches f_1 to ℓ_1, f_2 to ℓ_2, and f_3 to the null. The reason f_3 is matched to the null is that the cardinal monotonicity condition forces ℓ_2 to permanently reject f_3 when f_2 and f_3 propose, as both families would be rejected if f_1 also proposed. However, in the TKDAC algorithm, f_1 clinches ℓ_1 and goes to the bottom of ℓ_2's priority; as a result, f_2 and f_3 clinch ℓ_2. Therefore, clinching allows identifying that f_1 will not propose to ℓ_2; hence, ℓ_2 can accept both f_2 and f_3 as the violation of the cardinal monotonicity condition created will not materialize.

TKDAC – Round 1

ℓ_1	$(4,2)$	θ
f_1	$(2,1)$	∞
f_2	$(1,0)$	∞
f_3	$(1,0)$	∞
$\boxed{f_4}$	$(2,0)$	2
$\boxed{f_5}$	$(1,1)$	2
$\boxed{\cancel{f_6}}$	$(1,0)$	2
f_7	$(1,0)$	0

ℓ_2	$(4,1)$	θ
f_5	$(1,1)$	∞
$\boxed{f_1}$	$(2,1)$	1
f_6	$(1,0)$	∞
$\boxed{\cancel{f_3}}$	$(1,0)$	1
f_4	$(2,0)$	0
f_2	$(1,0)$	0
f_7	$(1,0)$	0

ℓ_3	$(3,1)$	θ
f_4	$(2,0)$	∞
f_6	$(1,0)$	∞
f_1	$(2,1)$	1
$\vdots$		

ℓ_4	$(1,0)$	θ
$\boxed{f_7}$	$(1,0)$	∞
$\boxed{\cancel{f_2}}$	$(1,0)$	0
$\vdots$		

TKDAC – Round 2

ℓ_1	$(4,2)$	θ
f_1	$(2,1)$	∞
$\boxed{f_2}$	$(1,0)$	∞
$\boxed{f_3}$	$(1,0)$	∞
$\boxed{f_4}$	$(2,0)$	3
$\boxed{\cancel{f_5}}$	$(1,1)$	0
f_6	$(1,0)$	0
f_7	$(1,0)$	0

ℓ_2	$(4,1)$	θ
f_5	$(1,1)$	∞
$\boxed{f_1}$	$(2,1)$	1
$\boxed{f_6}$	$(1,0)$	∞
f_3	$(1,0)$	1
f_4	$(2,0)$	0
f_2	$(1,0)$	0
f_7	$(1,0)$	0

ℓ_3	$(3,1)$	θ
f_4	$(2,0)$	∞
f_6	$(1,0)$	∞
f_1	$(2,1)$	1
$\vdots$		

ℓ_4	$(1,0)$	θ
$\boxed{f_7}$	$(1,0)$	∞
f_2	$(1,0)$	0
$\vdots$		

TKDAC – Round 3

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

ℓ_2	$(4,1)$	θ
$\boxed{f_5}$	$(1,1)$	∞
$\boxed{\cancel{f_1}}$	$(2,1)$	0
$\boxed{f_6}$	$(1,0)$	∞
f_3	$(1,0)$	0
f_4	$(2,0)$	0
f_2	$(1,0)$	0
f_7	$(1,0)$	0

ℓ_3	$(3,1)$	θ
f_4	$(2,0)$	∞
f_6	$(1,0)$	∞
f_1	$(2,1)$	2
$\vdots$		

ℓ_4	$(1,0)$	θ
$\boxed{f_7}$	$(1,0)$	∞
f_2	$(1,0)$	0
$\vdots$		

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

Preference type	Type 1	Type 2	Type 3	Type 4
Largest-to-smallest family	5.9 (1.8)	50.0 (15.2)	49.5 (15.0)	23.7 (7.2)
Smallest-to-largest family	6.6 (2.0)	44.6 (13.5)	46.0 (14.0)	23.3 (7.1)

Table F.1. Number (percentage) of families made better off for different orderings in the Rejection Stage of the KTTCE (one-dimensional constraints). Averages over 100 simulation rounds. Numbers are rounded to 1 d.p.

Preference type	Type 1	Type 2	Type 3	Type 4
Families made better off	6.8 (2.1%)	10.7 (3.3%)	11.7 (3.6%)	9.7 (3.0%)

Table F.2. Number (fraction) of families made better off by the KTTCE mechanism (three-dimensionsional constraints). Averages over 100 simulation rounds.

Appendix F. Further Simulation Results

F.1. **Effect of the order of families in the Rejection Stage in the KTTCE algorithm.** 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 (same-size families continue to be ordered randomly). The results of the two alternative family orderings in the Rejection Stage are displayed in Table F.1. We see that the largest-to-smallest ordering is at least as efficient as the smallest-to-largest ordering: the differences are pronounced for Type 2 and Type 3 preferences which generate the most improvements and are not significant for Type 1 and Type 4 preferences. 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.

Preference type	Type 1	Type 2	Type 3	Type 4
	Interference violations			
KTTCE	18888 (17.4%)	18225 (16.8%)	18378 (17.0%)	18617 (17.2%)
KTTC	603 (0.6%)	1302 (1.2%)	1306 (1.2%)	4390 (4.1%)
KDA	0	0	0	0
TKDA	0	0	0	0
	Average priority rank			
KTTCE	108	108	108	108
KTTC	36	97	96	78
KDA	25	27	27	26
TKDA	14	13	14	14
	Number of matched families			
KTTCE	315 (95.8%)	315 (95.8%)	315 (95.8%)	315 (95.8%)
KTTC	303 (92.1%)	307 (93.4%)	308 (93.4%)	302 (91.9%)
KDA	286 (86.7%)	287 (87.1%)	289 (86.1%)	286 (86.9%)
TKDA	234 (71.1%)	230 (70.0%)	231 (70.2%)	232 (70.7%)
	Fraction of unfilled capacity			
KTTCE	9.5%	9.4%	9.4%	9.4%
KTTC	11.6%	12.3%	12.3%	13.1%
KDA	17.8%	15.8%	15.9%	16.8%
TKDA	26.6%	27.3%	27.5%	27.1%

TABLE F.3. Outcomes of KTTCE, KTTC, KDA, and TKDA algorithms (one-dimensional constraints). Averages over 100 simulation rounds (and, for average priority rank, over all localities). Numbers (percentages) are rounded to the nearest integer (to 1 d.p.).

F.2. **Performance of KTTC, KTTCE, KDA, and TKDA algorithms in the three-dimensional setting.** Table F.2 summarizes the performance of the KTTCE mechanism in the three-dimensional environment. Compared to the one-dimensional case, fewer families are made better off on average. This finding is not surprising since the families' needs can differ across three dimensions, which makes feasible trades harder to find. However, despite our result that no mechanism is guaranteed to improve upon an endowment when there are multiple services (Theorem 2), we see that the KTTCE mechanism continues to find a significant number of improvements.[25]

Table F.3 and Figure F.1 display the results for the KTTC, KDA, and TKDA algorithms in the three-dimensional environment. The results are consistent with the one-dimensional environment (Table 6) across preference types and mechanisms. The second panel of Table F.3 shows that the average rank of matched families across different algorithms in the

[25]Sizes are not monotonic because a family can have more adults but fewer children than another family.

three-dimensional setting is similar to that of the one-dimensional setting. Figure F.1 shows that the ranking of the mechanisms by efficiency in the one-dimensional setting (Figure 2) is preserved in the three-dimensional setting. Comparing Table 6 and Table F.3, we see that our mechanisms tend not to perform as well in the three-dimensional environment as in the one-dimensional environment: the KTTC mechanism generates more interference violations and all three mechanisms are less efficient. These findings are intuitive in the sense that having three dimensions in effect 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 interference violations; moreover, in the KDA and TKDA mechanism, localities may have to reject more families in order to prevent interference violations.

Finally, we assess the efficiency gain associated with using our solution concept of interference-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, there is no impact on the outcome of either algorithm (see Proposition C.3 and Corollary C.1). However, in the presence of multidimensional constraints, Figure F.2 shows that the impact on efficiency of using interference-freeness over envy-freeness is substantial both for the KDA and the TKDA algorithms.

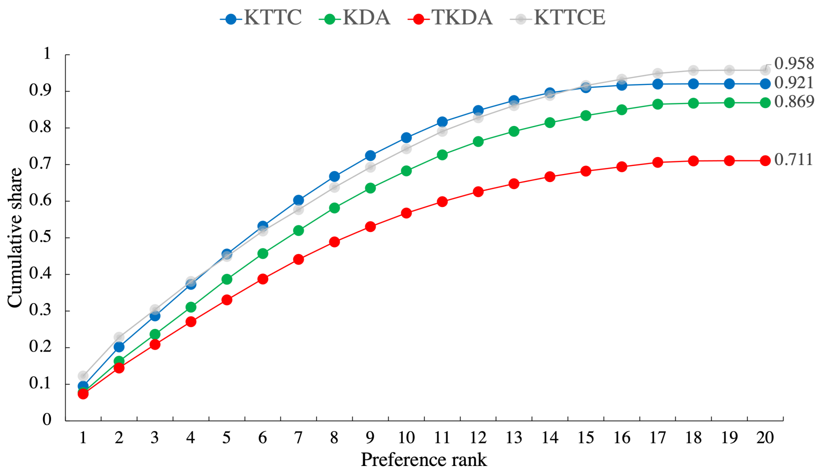

(A) Type 1

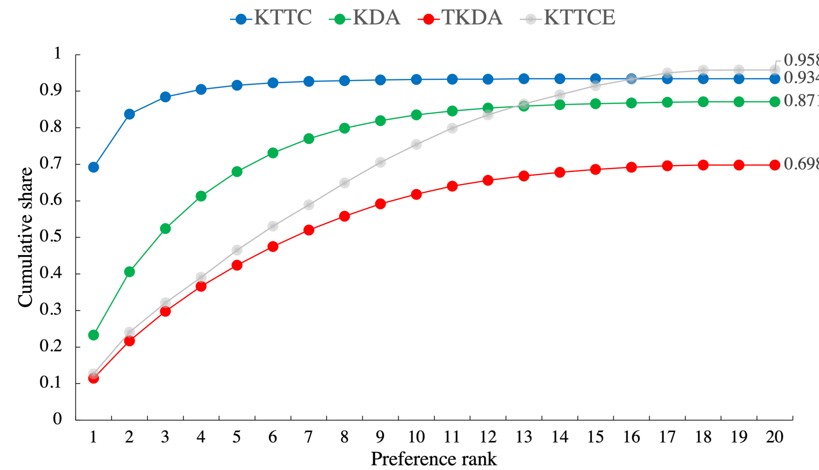

(B) Type 2

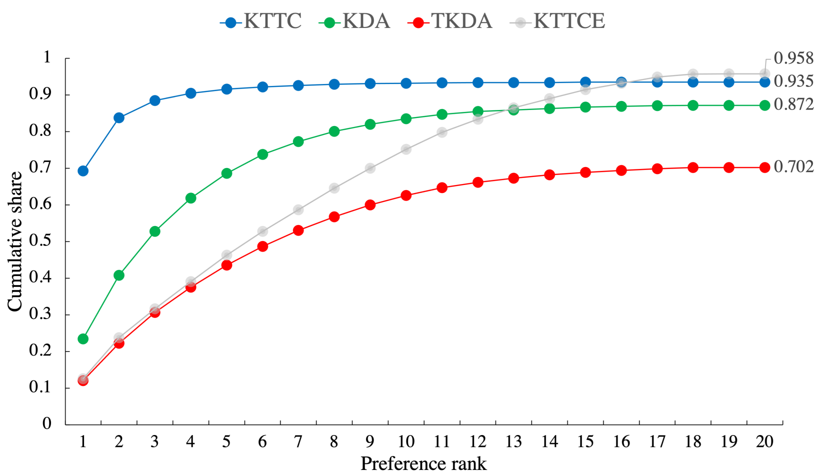

(C) Type 3

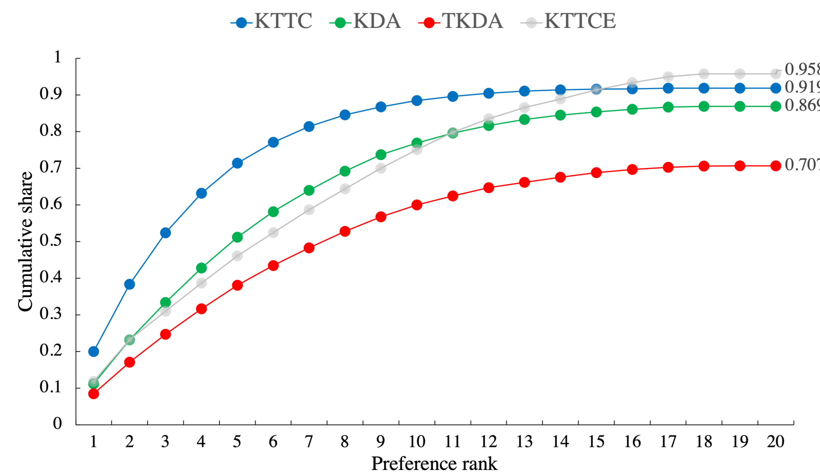

(D) Type 4

FIGURE F.1. Preference distributions for outcomes under different matching mechanisms (three-dimensionsional constraints). Labeled numbers: fractions of matched families.

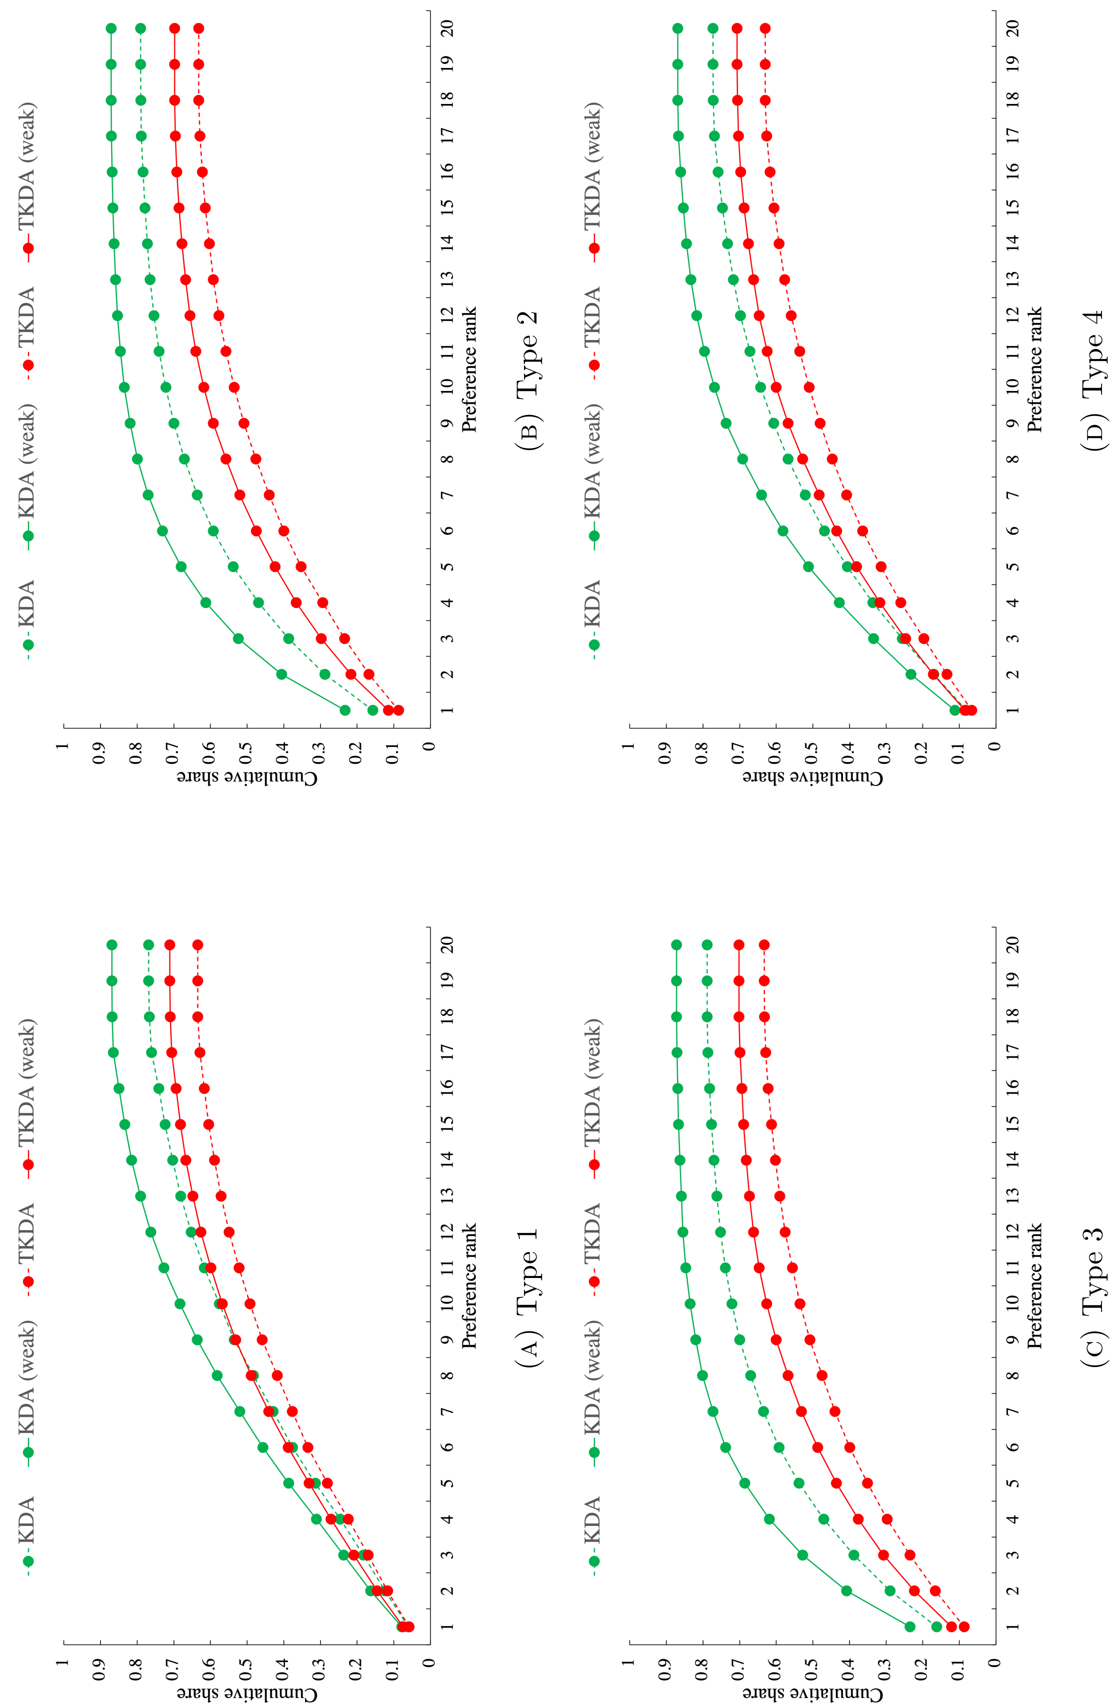

FIGURE F.2. Preference distributions for KDA and TKDA algorithms when using accommodation vs. weak accommodation (three-dimensional constraints).

Appendix G. Relationships to Prior Models

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

- School choice (Abdulkadiroğlu and Sönmez, 2003): Each 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 $f \in F$ having size $\nu_d^f = 1$.
- 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 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 an 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 $\nu^f = (1, 0, 0, \ldots, 1, \ldots, 0, 0)$ 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 $\{0, 1\}^3$.

- 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.[26]

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.

References

Abdulkadiroğlu, Atila, "College Admissions with Affirmative Action," *International Journal of Game Theory*, 2005, *33* (4), 535–549.

______ **and Tayfun Sönmez**, "School Choice: A Mechanism Design Approach," *American Economic Review*, 2003, *93* (3), 729–747.

Aygün, Orhan and Inácio Bó, "College Admission with Multidimensional Privileges: The Brazilian Affirmative Action Case," *American Economic Journal: Microeconomics*, forthcoming 2020.

Biró, Péter and Eric McDermid, "Matching with sizes (or scheduling with processing set restrictions)," *Discrete Applied Mathematics*, 2014, *164*, 61–67.

Biró, Péter, Tamás Fleiner, Robert W. Irving, and David F. Manlove, "The College Admissions problem with lower and common quotas," *Theoretical Computer Science*, 2010, *411* (34), 3136 – 3153.

Cechlárová, Katarína, Tamás Fleiner, David F Manlove, Iain McBride, and Eva Potpinková, "Modelling practical placement of trainee teachers to schools," *Central European Journal of Operations Research*, 2015, *23* (3), 547–562.

Delacrétaz, David, "Stability in Matching Markets with Sizes," Mimeo 2019.

Hafalir, Isa Emin, M. Bumin Yenmez, and Muhammed Ali Yildirim, "Effective Affirmative Action in School Choice," *Theoretical Economics*, 2013, *8* (2), 325–363.

[26]This 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.

Hatfield, John William and Paul Milgrom, "Matching with Contracts," *American Economic Review*, 2005, *95* (4), 913–935.

Kamada, Yuichiro and Fuhito Kojima, "Fair Matching Under Constraints: Theory and Applications," *Review of Economic Studies*, forthcoming.

Kojima, Fuhito, "School Choice: Impossibilities for Affirmative Action," *Games and Economic Behavior*, 2012, *75* (2), 685–693.

McDermid, Eric J and David F Manlove, "Keeping partners together: algorithmic results for the hospitals/residents problem with couples," *Journal of Combinatorial Optimization*, 2010, *19* (3), 279–303.

Roth, Alvin E., "Stability and polarization of interests in job matching," *Econometrica*, 1984, *52* (1), 47–57.

Westkamp, Alexander, "An analysis of the German university admissions system," *Economic Theory*, 2013, *53* (3), 561–589.