The Role of Priorities in Assigning Indivisible Objects: A Characterization of Top Trading Cycles

Atila Abdulkadiroglu and Yeon-Koo Che

Duke University and Columbia University

This version: November 2010
Preliminary

Abstract

The Top Trading Cycles mechanism emerges as a desirable solution in various market design applications. Yet recommendations are made without any rigorous foundation for the role priorities play in assignment. We explain that role by *respecting top priorities,* which generalizes the idea that if an agent is assigned an object, then the agent that is top-ranked by that object has to be assigned a weakly better object. We show that, given a priority structure, a mechanism is Pareto efficient, strategy-proof and recursively respects top priorities if and only if it is the Top Trading Cycles mechanism defined by the priority structure.

1 Introduction

Various notable market design applications involve allocation of heterogenous indivisible objects without monetary transfers, such as assigning pupils to public schools in a school choice program and rematching kidney patients with donors when patients have donors with incompatible kidneys. A common feature of such problems is that, objects usually rank individuals in a priority order, which resembles individuals' preference orderings of the objects. *Top Trading Cycles* (TTC) and its variants, the roots of which can be traced back to Gale's celebrated top trading cycles algorithm, emerges as a desirable solution to incorporate such priorities in the allocation process.

A typical problem consists of a finite set of individuals, say students, and a set of objects with finite capacities, say schools. Students rank schools in strict preference order and schools rank students in priority order. For example, a child who lives within a certain distance from a school may get neighborhood priority for that school and any tie in priorities may be broken by a fair lottery.

Given strict preference lists of students and strict priority lists of schools, TTC assigns students to schools via the following algorithm (Abdulkadiroğlu and Sönmez 2003): In each round, every student points to the school she prefers most among the remaining schools, and every remaining school points to the student that has the highest priority at that school among all remaining students. A cycle is an ordered list of schools and students $\{o_1, i_1, ..., o_K, i_K\}$ such that school o_k points to student i_k and student i_k points to school o_{k+1} and $o_{K+1} \equiv o_1$. When such a cycle exists, student i_k is assigned to school o_{k+1}, the capacity of school o_{k+1} is decreased by one, the students in the cycle and schools with no more capacity are removed; the process is repeated with the remaining students and schools.

The resulting assignment is *Pareto efficient*, that is, there is no alternative assignment that improves a student's assignment without harming others'. Furthermore, TTC is *strategy-proof*, i.e. it makes truthful reporting of preferences over schools a dominant strategy for every student in the induced preference revelation game.

However many other mechanisms meet these two requirements. For example, given an ordering of students, a corresponding *serial dictatorship* mechanism determines the assignment as follows: Each student is assigned in the given order to her most preferred school among the remaining ones. Any serial dictatorship mechanism is Pareto efficient and strategy-proof as well.

Given the richness of Pareto-efficient and strategy-proof mechanisms, then a natural question arises: Why TTC? In other words, what is the additional property that uniquely pins down TTC among all efficient and strategy-proof mechanisms? Despite its prevalence in market design applications, this question remains open. Without a satisfactory answer, any recommendation of TTC in such applications would not be well-grounded. We fill in this gap by offering a new characterization of TTC.

The most related work to ours is Pápai (2000). In a model in which each object has unit capacity, Pápai characterizes a wide class of mechanisms, hierarchical exchange rules, by Pareto efficiency, group strategy proofness, which rules out beneficial preference manipulation by groups of individuals, and reallocation-proofness, which rules out manipulation by two individuals via misrepresenting preferences and swapping objects ex post. TTC is a hierarchical exchange rule defined by the priority lists of schools. Yet, none of the properties above makes any explicit reference to any specific priority structure. In fact, one of them is about efficiency and the other two are about avoiding sophisticated manipulation schemes by individuals. Therefore, Pápai's result does not address the question of why those specific priorities are used in TTC.[1]

Any normative explanation to the use of specific priorities in TTC should involve an axiom that makes a reference to the priorities. Our explanation rests on the idea of respecting top priorities in the following sense: If a student is assigned a school, than the top ranked student at that school should be satisfied first, that is she has to be assigned a school that she weakly prefers to that school, a property that we call *respecting top priorities.* This requirement is very weak that it is satisfied even by the null matching that does not assign any student to any school. TTC satisfies that criterion, but a serial dictatorship mechanism does not, which is also a hierarchical exchange rule.

However, we would like to attribute a role to the full priority structure, not only to the top priorities. The requirement of "respecting top priorities" generalizes intuitively for the full priority structure. Suppose that a student i is assigned school s_1, and the top ranked student i_1 at s_1 has to be assigned a school s_2 that she weakly prefers to s_1; then the top ranked student i_2 at s_2 has to be assigned a school s_3 that she weakly prefers to s_2; and so on. By finiteness we may reach back to i_1 during that process. Then $\{i_k\}_{k=1,\dots,K}$ forms a *top priority students for* i, and these top priority students have to be satisfied for i to be assigned s_1. A matching that respects top priorities satisfies those students. Once those students are removed with their assignments, the matching continues to respect top priorities in the reduced problem. If that property is satisfied for all such reduced problems recursively, we say

[1] In a similar vein, Pycia and Ünver (2010) introduce and characterize trading cycles with brokers and owners by Pareto efficiency and group strategy-proofness. Also Kojima and Manea (2010) provide a characterization for Gale-Shapley and Kojima and Unver (2010) provide a characterization for the Boston school choice mechanism, both via a monotonicity condition on preferences.

that the matching *rcursively respects top priorities.*

Our main result states that a mechanism is Pareto efficient, strategy proof and recursively respects top priorities if and only if it is the TTC defined by the priority structure. The intuition behind that characterization result is as follows: Consider any Pareto efficient and strategy-proof mechanism φ that respects priorities. Consider the schools that are assigned in the first round of TTC. If φ is Pareto efficient, then it has to assign some student to some of those schools (note that the null assignment always respects priorities). If it ever assigns a student to one of these schools, then the students that are assigned in the first round of TTC have to be satisfied. Pareto efficiency and strategy-proofness together imply that they are assigned their first choices, which is their TTC assignment. Once these students are assigned (or satisfied) in a strategy proof and efficient way, it is easy to generalize this argument in the reduced problem in which those students are removed with their assignments, so that φ coincides with TTC.

Such foundation is important in policy making. A common interpretation of TTC is that it effectively allows students to trade in their priorities for better schools of their choice. However that interpretation may prove difficult in making a case for TTC. For instance, sibling priority is usually granted on the belief that assigning siblings to the same schools benefits the siblings via spillover effects and sharing experiences and their parents via solving transportation and coordination problems. From the point of view of a policy maker, trading in sibling priority for a better choice may be difficult to justify since the sibling priority may have been instituted for encouraging siblings to go to the same school. In contrast, our characterization states that, when a school district has Pareto efficiency, strategy-proofness and respecting priorities as three policy goals to meet, the unique mechanism that meets those criteria is the TTC defined by the given priority structure. It is worth noting that there is no reference to any trading of priorities in the three stated goals, therefore TTC is justified not by allowing students to trade in their priorities but by three policy goals none of which require trading of priorities.

Our result provides a clear answer to the question of what role priorities play in the allocation of indivisible objects. Another intuitive role that can be attributed to priorities is a monotonicity relation between priorities and the assignment. Namely, a mechanism *respects improvements in priorities* if whenever a student's standing in priorities improves, her assignment weakly improves. Note that this notion does not make any reference to any specific priority structure, so it cannot be used to pin down a single priority structure to define TTC. However, TTC satisfies this requirement. Therefore, as a corollary of our main result, Pareto efficiency, strategy-proofness and recursive individual rationality also implies respecting improvements in priorities.

We formalize our arguments in the following sections.

2 Model

A problem consists of a finite set of agents $I = \{1, ..., n\}$ and a finite set of objects $O = \{a, b, c, ...\}$. To simplify the exposition, we will assume that each object has a single copy, but the arguments can be easily generalized if objects come in multiple copies, such as schools with multiple seats in school choice. An agent can consume at most one object and an object can be consumed by at most one agent. Each agent $i \in I$ has a complete, irreflexive and transitive binary preference relation P_i over $O \cup \{\varnothing\}$ and $\varnothing$ represents consuming nothing. aP_ib means that i prefers a to b. Each object $a \in O$ ranks agents by a complete, irreflexive and transitive binary priority relation $\succ_a$ over A. $i \succ_a j$ means that i has higher priority at a than j.

Let $P = (P_i)_{i\in I}$, $\succ = (\succ_a)_{a\in O}$, $P_{-I'} = ((P_j)_{j\in I-I'})$ and $\succ_{-O'} = (\succ_b)_{b\in O-O'}$. We fix I and O and refer to a problem by $(P, \succ)$.

For $i \in I$ let R_i be the symmetric extension of P_i, that is, for all $a, b \in O \cup \{\varnothing\}$, if aP_ib then aR_ib, and if $a = b$ then aR_ib and bR_ia. Let the indifference relation I_i denote the symmetric part of R_i. Define $\succsim$ similarly.

A **matching** of agents to objects is a function $\mu : A \longrightarrow O$ such that $\mu(i) \subset O$, $|\mu(i)| \leq 1$ for all $i \in I$.

A matching μ (Pareto) **dominates** another matching ν if $\mu(i)R_i\nu(i)$ for all $i \in I$ and $\mu(i)P_i\nu(i)$ for some $i \in I$.

A matching is **Pareto efficient** if it is not dominated by another matching.

A (deterministic) **mechanism** selects a matching for every problem. A mechanism is efficient if it selects an efficient matching for every problem. If φ is a mechanism, let $\varphi(P; \succ)$ denote the matching selected by φ. A mechanism φ is **strategy-proof** if reporting true preferences is a dominant strategy for every agent in the preference revelation game induced by φ, that is

$$\varphi(P; \succ)(i)R_i\varphi(P'_i, P_{-i}; \succ)(i) \tag{1}$$

for all P, $\succ$, $i \in I$ and P_i'.

3 Top Trading Cycles (TTC)

Given a problem $(I, O, P, \succ)$, TTC finds the matching via the following algorithm:

- In the first round of the algorithm, all students and schools are available. In every round of the algorithm,
- Every available object points to its highest priority agent among all available agents. Every agent points to her most preferred object among all available objects.
- A *cycle* $c = \{o_k, i_k\}_{k=1,...,K}$ is an ordered list of objects and agents such that o_k points to i_k and i_k points to o_{k+1} for every $k = 1, ..., K$, where $o_{K+1} = o_1$.
- For every cycle $c = \{o_k, i_k\}_{k=1,...,K}$, match each agent with the object she points to in that cycle and remove the agent and the object. In that case, we say that i_k *trades* o_k *for* o_{k+1}.
- Repeat the algorithm in the next round until no more agents are matched.

4 Charaterization: Respecting Priorities

Definition 1 *Given $\succ$, a matching μ* ***respects top priorities*** *if for every student i, the top ranked student j at $\mu(i)$ prefers $\mu(j)$ to $\mu(i)$ (we also say that j is satisfied).*

Note that this is a very weak requirement. For instance, the null matching that leaves all students unmatched trivially respects top priorities. It puts a restriction on the matching when a student is assigned a school at which she is not ranked highest. In that case, the top ranked student at that school should receive something better. This is implied by the following intuitive requirement.

Definition 2 *Given $\succ$, a matching μ is* ***individually rational for top students*** *if every student that is ranked highest at a school is matched with an alternative that she weakly prefers to that school.*

Respecting top priorities is weaker than individual rationality for top students. For instance, the null matching respects top priorities, but it is not individually rational for top ranked students when a student prefers some school to being unassigned. In a Shapley-Scarf economy, every agent owns a house. This can be modelled as a special case of our model in which every agent is ranked highest at the house she owns. In that economy, individual rationality for top students coincides with with Ma's individual rationality.

Consider a matching μ that respects top priorities and a student i who is assigned $\mu(i) \in O$. Let i' be the top ranked student at $\mu(i)$. If i' is not assigned, then i' prefers being unassigned to being assigned $\mu(i)$. If i' is assigned some school, i.e. $\mu(i') \in O$, then consider the top ranked student i'' at $\mu(i')$. If i'' is assigned some school, i.e. $\mu(i'') \in O$, then consider the top ranked student i''' at $\mu(i'')$. Eventually we obtain a set of students $\{i_k\}_{k=1,...,K}$ such that i_{k+1} is the highest ranked student at $\mu(i_k)$ where $i_{K+1} \equiv i_1$. For μ to respect top priorities, it must be the case that $\mu(i_{k+1}) \succsim_{i_{k+1}} \mu(i_k)$ for all $k = 1, ..., K$ where $i_{K+1} \equiv i_1$. That is, respecting top priorities leads to the following natural notion:

Definition 3 *An ordered list of students $\{i_k\}_{k=1,...,K}$, $K \geq 1$, forms a* ***top priority group*** *for a student i if i_1 is the top ranked student at $\mu(i)$ and either $K = 1$ and i_1 is unassigned or i_{k+1} is top ranked at $\mu(i_k)$ for every $k = 1, ..., K$, where $i_{K+1} \equiv i_1$.*

If a matching assigns some students to some schools, then a top priority group (for some student) exists by finiteness. If μ respects top priorities, then the students in the top priority group of every student are satisfied.

Given $\succ$, a matching μ and a set of students I'', define a subproblem by removing the students in I'' with their assignment. That is, the set of students in the subproblem is $I' = I \backslash I''$, the set of objects is $O' = O \backslash \cup_{i \in I''} \mu(i)$, preferences of I' and priorities at O' are obtained by projection of $\succ$ on I' and O'.

In this section, we will consider certain subproblems: Given $\succ$, a matching μ and a student i and a top priority group $\{i_k\}_{k=1,...,K}$, define a subproblem by removing the students in $\{i_k\}_{k=1,...,K}$ with their assignment. That is, the set of students in the subproblem is $I' = I \backslash \{i_k\}_{k=1,\ldots,K}$, the set of objects is $O' = O \backslash \cup_{i \in \{i_k\}_{k=1,\ldots,K}} \mu(i)$, preferences of I' and priorities at O' are obtained by projection of $\succ$ on I' and O'.

The following recursive definition of respecting priorities generalizes the notion of respecting top priorities to full priority structure.

Definition 4 *Given $\succ$, a matching μ* ***recursively respect top priorities*** *if it respects top priorities, and its projection in the subproblem obtained by removing students and their assigned objects in any top priority group recursively respects top priorities.*

Theorem 5 *A mechanism φ is Pareto efficient, strategy-proof and recursively respects priorities if and only if $\varphi(P, \succ) = TTC(P, \succ)$ for all $(P, \succ)$.*

We defer the proof to the appendix.

Note that the null matching mechanism is strategy-proof and trivially respects priorities, but it is not Pareto efficient. A serial dictatorship mechanism is Pareto efficient and strategy-proof, but it does not respect priorities. For a mechanism that is PE, respect priorities, but not SP, modify TTC as follows: Pick a student i. If i is not top ranked at her top two most preferred schools, and if her most preferred school is also somebody else's most preferred school, she points to her second choice in the first round of TTC, otherwise it is the usual TTC.

5 The Connection with Stability

A pair (i, o) pair **blocks** a matching μ if $\mu(o) \in I$, $o \succ_i \mu(i)$ and $i \succ_o \mu(o)$. A matching is **stable** if it is not blocked.

The standard notion of blocking pair does not require $\mu(o) \in I$ and accordingly (i, o) may block μ when o is not assigned under μ. But that kind of blocking is more about efficiency in a one-sided matching framework so it is taken care of by the efficiency requirement.

Note that stability can be defined via the notion of respecting top priorities as follows:

Theorem 6 *A matching is stable if and only if and only if its projection in any subproblem obtained by removing any subset of students and their assigned objects respects top priorities.*

One of the difficulties in the matching literature has been to understand the difference and similarities between TTC and Gale and Shapley's student optimal stable matching mechanism. Our axiom of "respecting top priorities" provides a unified framework in which we can compare the two mechanisms. Accordingly, both mechanisms are charaterized by strategy proofness and Pareto optimality sucbject to respecting top priorrities: While Gale-Shapley checks for respecting top priorities in all subproblems, TTC checks for a limited set of subproblems. Therefore, while TTC obtains full Pareto efficiency, Gale-Shapley obtains only efficiency with the set of stable matching.

6 The Shapley-Scarf Economy

In a housing market problem (Shapley and Scarf 1974) each individual owns a house which she would like to exchange for another one she prefers more. The unique core of this market is found via Gale's top trading cycles algorithm (Postlewaite and Roth 1977), in which each agent points to the owner of his most preferred house among the remaining houses. An allocation is said to be individually rational for a housing market if every agent is assigned a house that she weakly prefers to her initial endowment. Ma (1994) shows that Gale's top trading cycles is characterized by Pareto efficiency, strategy proofness and individual rationality. It is in that sense that our generalized axiom provides a foundation for TTC.

Although Shapley and Scarf motivates the problem as an exchange economy for houses in which every agent owns a single house, the model finds applications beyond that market. For instance, the houses may represent campus housing for college students and owning a house may mean being an existing tenant for that house. Individual ratinality would not have any substantive interpretation in that environment.

In that environment, we say that a matching μ **does not violate the ownership structure** if for every i, the owner of $\mu(i)$ prefers $\mu(j)$ to $\mu(i)$. Note that this is much weaker than individual rationality as the null matching satisfies it. The the following result is a corrollary to our Theorem:

Theorem 7 *A mechanism is strategy-proof, Pareto efficient and does not violate the ownership structure if and only if it is Gale's TTC.*

Note that we do not need individual rationality for this characterization.

7 Monotonicity

Another intuitive role that can be attributed to priorities is a monotonicity relation between priorities and the assignment. Namely, if whenever a student's standing in priorities improves, her assignment is expected to improve. We make this formal below.

Definition 8 *$\succ'$ is an improvement in priorities for $i \in I$ if*

$$\begin{aligned}
&\succ' \text{ is not equivalent to } \succ, \\
&i \succ j \Rightarrow i \succ' j \text{ and} \\
&\forall j, k \in I - \{i\} : j \succ' k \Leftrightarrow j \succ k
\end{aligned}$$

Definition 9 *A mechanism φ is respects improvements in priorities if for all $(P, \succ)$, $i \in I$*

(i) if $\succ'$ is an improvement for i, then $\varphi(P, \succ')(i) R_i \varphi(P, \succ)(i)$; and

(ii) if $\varphi(P, \succ)(i)$ is not i's first choice, then there exists an improvement $\succ'$ for i such that $\varphi(P, \succ')(i) P_i \varphi(P, \succ)(i)$.

Since this notion does not make any reference to any specific priority structure, it does not pin down a TTC with a particular priority structure from the set of Pareto efficient and strategy-proof mechanisms. However, TTC satisfies this requirement. Therefore, as a corollary of our main result, Pareto efficiency, strategy-proofness and recursive individual rationality also implies respecting improvements in priorities.

Corollary 10 *Pareto efficiency, strategy-proof and respecting priorities implies respecting improvements in priorities.*

8 Extensions and Discussion

Priorities and priority-based mechanisms play an essential role in the allocation of indivisible objects when monetary transfers are not allowed. The role of priorities in Gale-Shapley's celebrated student proposing deferred acceptance mechanism (Gale-Shapley) is well understood. Gale-Shapley is characterized as the student optimal stable stable matching mechanism. Likewise, the Boston mechanism is the student optimal mechanism which produces a stable matching according to the preference-adjusted priorities in which a student who rank a school higher in her choice list has higher priority than a student who ranks it lower and they are ranked according to the original priority order otherwise. Our result complements the picture by explaining the role priorities play in TTC.

To simplify the exposition of the ideas, we have assumed that each object comes in single copy When objects have multiple copies, such as schools in school choice, an object is removed in the definition of recursive individual rationality when all of its copies are removed in the recursive process. With this modification to the definition of the axiom, the main results follows without any change.

A second assumption is that objects rank individuals in strict priority order. The main result follows directly after the breaking of ties potentially via some lottery but not before. To see this, suppose that there is only one school with one seat and two students with equal priority at the school. Since no mechanism can give both the students the single available seat, no mechanism can respect top priorities for all students ranked highest.

When ties at school priorities are broken randomly, an interesting monotonicity relation between priorities and the random TTC allocation emerges. Namely, consider an improvement in priorities for a student. Then the student's random TTC allocation under the improved priority structure first order stochastically dominates her random TTC allocation under the original priority structure. This follows from the fact that, for any tie breaking, her assignment weakly improves under the improved priority structure. A characterization of random TTC remains an open question.

References

[1] **Abdulkadiroğlu, Atia, and Tayfun Sönmez.** 2003. "School Choice: A Mechanism Design Approach." *American Economic Review*, 93, 729-747.

[2] **Kojima, Fuhito and Mihai Mane** 1994. "Axioms for Deferred Acceptance." *Econometrica*, forthcoming

[3] **Kojima, Fuhito and Utku Ünver** 1994. "The Boston School-Choice Mechanism." mimeo

[4] **Ma, Jinpeng** 1994. “Strategy-Proofness and the Strict Core in a Market with Indivisibilities.” *International Journal of Game Theory*, 23, 75-83

[5] **Pápai, Szilvia** 2000. “Strategyproof Assignment by Hierarchical Exchange.” *Econometrica*, Vol. 68, No. 6 (Nov., 2000), pp. 1403-1433.

[6] **Pycia, Marek and Utku Ünver** 2010. “Incentive Compatible Allocation and Exchange of Discrete Resources.” *mimeo*

[7] **Postlewaite, A. and A.E. Roth** 19tt. “Weak versus Strong Domination in a Market with Indivisibile Goods.” *Journal of Mathematical Economics*, 4, 131-137

[8] **Shapley, LS and H. Scarf** 1974. “On cores and Indivisibility.” *Journal of Mathematical Economics*, 1, 23-37

A Proof of Theorem 1

TTC is Pareto efficient and strategy-proof (Abdulkadiroglu and Sonmez 2003).

Consider any student i who is top ranked at some object o. Object o remains and points to i until i is assigned in TTC. If all objects that i prefers to o' are assigned to other agents at some round of TTC, then i points to o, $\{o, i\}$ forms a cycle and i is assigned o. Otherwise i is assigned an object she prefers to o. So $TTC(P, \succ)(i) \succsim_i o$ for all such i. In other words, $TTC(P, \succ)$ respects top priorities.

Consider a sequence of student groups $\{i_1^t, ..., i_{K_t}^t\}_{t=1,...,T}$ such that $\{i_1^1, ..., i_{K_1}^1\}$ is a top priority group in the orginal problem and for each $t' = 2, ..., T$ the set $\{i_1^{t'}, ..., i_{K_{t'}}^{t'}\}$ is a top priority group in the reduced problem when students in $\{i_1^t, ..., i_{K_t}^t\}_{t=1,...,t'-1}$ and their assignments are removed. We will prove the statement by induction on t'.

For the induction step, suppose that $TTC(P, \succ)$ respects top priorities in the sub problem when agents in $\{i_1^t, ..., i_{K_t}^t\}_{t=1,...,t'-1}$ and their assignments are removed, where $t' \geq 1$. We have proved the induction hypothesis for $t' = 1$. We will prove that $TTC(P, \succ)$ respects top priorities in the new sub problem when agents in $\{i_1^{t'}, ..., i_{K_{t'}}^{t'}\}$ and their assignments are also removed.

Note that the TTC algorithm can be restated as follows: Let every remaining agent point to her most preferred object among all the remaining objects and every remaining object points to its top ranked agent among all remaing agents. Instead of implement all the cycles, pick one cycle, implement the cycle, reduce the problem by removing the agents and their assignments in the cycle; repeat this process with the reduced problem. When agents in a cycle are removed with their assignments, the projection of $TTC(P, \succ)$ on the sub problem coincides with the TTC matching of the subproblem.

Any top priority group, and in particular $\{i_1^1, ..., i_{K_1}^1\}$, corresponds to the set of students involved in a cycle in the first round of TTC. Consider the sub problem obtained by removing agents in $\{i_1^t, ..., i_{K_t}^t\}_{t=1,...,t'-1}$. Any top priority group in the subproblem, and in particular $\{i_1^{t'}, ..., i_{K_{t'}}^{t'}\}$, corresponds to the set of students involved in a cycle in the first round of TTC in that subproblem. So by removing $\{i_1^{t'}, ..., i_{K_{t'}}^{t'}\}$ and their assignment, we obtain a subproblem for which the TTC matching coincides with the projection of $TTC(P, \succ)$ on the subproblem after the removal of agents in $\{i_1^t, ..., i_{K_t}^t\}_{t=1,...,t'}$ and their assignments. Consider any top ranked student in that subproblem. Since an object continues to point to its top ranked agent until the agent is assigned, every top ranked agent is guaranteed an assignment that she weakly prefers to the objects she is ranked highest in that subproblem. So $TTC(P, \succ)$ respects top priorities in that subproblem. This completes the induction step.

Next we will prove the other direction of the statement. Let φ be a Pareto efficient, strategy-proof and recursively respects top priorities. Let $\tilde{I}_k(P, \succ)$ be the set of agents who are matched in step k of $TTC(P, \succ)$.

Claim 1: $\varphi(P, \succ)(i) = TTC(P, \succ)(i)$ for all $i \in \tilde{I}_1(P, \succ)$.

Proof of Claim 1: Suppose to the contrary that $\varphi(P, \succ)(i) \neq TTC(P, \succ)(i)$ for some $i \in \tilde{I}_1(P, \succ)$.

Let $c = \{o_k, i_k\}_{k=1,\ldots,K}$ be the cycle in which i is matched with $TTC(P, \succ)(i)$ and $i = i_K$.

Note that every i_k trades o_k for o_{k+1}, which is her first choice. Consider the alternative preference relation $P'_{i_K} : o_1 o_K \varnothing$. By construction, the TTC matching remains the same, i.e.

$$TTC(P'_{i_K}, P_{-\{i_K\}}, \succ) = TTC(P, \succ).$$

Since o_1 is i_K's first choice and

$$\varphi(P, \succ)(i_K) \neq TTC(P, \succ)(i_K),$$

$o_1 P_{i_K} \varphi(P, \succ)(i_K)$. So by strategy-proofness, either $\varphi(P'_{i_K}, P_{-\{i_K\}}, \succ)(i_K) = \varnothing$ or $\varphi(P'_{i_K}, P_{-\{i_K\}}, \succ)(i_K) = o_K$. If $\varphi(P'_{i_K}, P_{-\{i_K\}}, \succ)(i_K) = \varnothing$, then respecting top priorities implies that no agent is assigned o_K since i_K is top-ranked by o_K. Then this contradicts with Pareto efficiency of φ since i_K prefers o_K to not being matched under P'_{i_K}. So $\varphi(P'_{i_K}, P_{-\{i_K\}}, \succ)(i_K) = o_K$.

Then

$$\varphi(P'_{i_K}, P_{-\{i_K\}}, \succ)(i_{K-1}) \neq TTC(P'_{i_K}, P_{-\{i_K\}}, \succ)(i_{K-1}) = o_K.$$

Now consider the alternative preference relation $P'_{i_{K-1}} : o_K o_{K-1} \varnothing$. By construction, the TTC matching remains the same, i.e.

$$TTC(P'_{i_{K-1}}, P'_{i_K}, P_{\{-i_K\}}, \succ) = TTC(P'_{i_K}, P_{-\{i_K\}}, \succ) = TTC(P, \succ).$$

Since o_K is i_{K-1}'s first choice and

$$\varphi(P'_{i_K}, P_{-\{i_K\}}, \succ)(i_{K-1}) \neq TTC(P'_{i_K}, P_{-\{i_K\}}, \succ)(i_{K-1}),$$

we obtain

$$o_K P_{i_{K-1}} \varphi(P'_{i_K}, P_{-\{i_K\}}, \succ)(i_{K-1}).$$

So by strategy-proofness, either $\varphi(P'_{i_{K-1}}, P'_{i_K}, P_{-\{i_{K-1},i_K\}}, \succ)(i_{K-1}) = \varnothing$ or $\varphi(P'_{i_{K-1}}, P'_{i_K}, P_{-\{i_{K-1},i_K\}}, \succ)(i_{K-1}) = o_{K-1}$. If $\varphi(P'_{i_{K-1}}, P'_{i_K}, P_{-\{i_{K-1},i_K\}}, \succ)(i_{K-1}) = \varnothing$, then respecting top priorities implies that no agent is assigned o_{K-1} since i_{K-1} is top-ranked by o_{K-1}. Then this contradicts with Pareto efficiency of φ since i_{K-1} prefers o_{K-1} to not being matched under $P'_{i_{K-1}}$. So $\varphi(P'_{i_{K-1}}, P'_{i_K}, P_{-\{i_{K-1},i_K\}}, \succ)(i_{K-1}) = o_{K-1}$.

Then

$$\varphi(P'_{i_{K-1}}, P'_{i_K}, P_{-\{i_{K-1},i_K\}}, \succ)(i_{K-1}) \neq TTC(P'_{i_{K-1}}, P'_{i_K}, P_{-\{i_{K-1},i_K\}}, \succ)(i_{K-2}) = o_{K-1}.$$

Repeating this argument recursively for every agent in the cycle c, we obtain that $\varphi(P'_c, P_{-c}, \succ)(i_1) = o_1$, where $P'_c = \{P'_{i_k}\}_{i_k \in c}$ and $P'_{i_k} : o_{k+1} o_k \varnothing$. Then $\varphi(P'_c, P_{-c}, \succ)(i_1) = o_1$ implies $\varphi(P'_c, P_{-c}, \succ)(i_K) \neq o_1$. Then if $o_K P'_{i_K} \varphi(P'_c, P_{-c}, \succ)(i_K)$, o_K is not asigned to any agent by φ since φ respects top priorities and o_K ranks i_K highest. But that contradicts with Pareto efficiency of φ since $o_K P'_{i_K} \varphi(P'_c, P_{-c}, \succ)(i_K)$. So $\varphi(P'_c, P_{-c}, \succ)(i_K) = o_K$, which implies $\varphi(P'_c, P_{-c}, \succ)(i_{K-1}) \neq o_K$. Repeated application of this argument proves that $\varphi(P'_c, P_{-c}, \succ)(i_k) = o_k$ for every i_k in the cycle. Then this contradicts with Pereto efficiency of φ because every agent in the cycle will be better off if every i_k is matched with o_{k+1} without changing the matching of agents in $I - c$.■

For any i, let o_i denote the object that agent i trades for her assignment in $TTC(P, \succ)$.

Now suppose to the contrary that $\varphi(P, \succ) \neq TTC(P, \succ)$. Then there must exist some i such that $TTC(P, \succ)(i) P_i \varphi(P, \succ)(i)$ and j such that $\varphi(P, \succ)(j) P_j TTC(P, \succ)(j)$, otherwise it would contradict with Pareto efficiency of φ and TTC. Furthermore $i \notin \tilde{I}_1(P, \succ)$ by Claim 1.

The proof is completed by repeated reduction of the problem. Pick any $\tilde{\imath} \in \tilde{I}_1(P, \succ)$ and consider the subproblem in which $\tilde{\imath}$'s top priority group $T(\tilde{\imath}; \mu, \succ)$ and their assignments are removed. The projection of μ respects top priorities in the reduced problem. Repeat this until every agent in $\tilde{I}_1(P, \succ)$ is removed. If i is assigned $TTC(P, \succ)(i)$ in the second round of TTC, then o_i has not been removed in the reduced problem since φ and TTC assignments are the same for all agents in $\tilde{I}_1(P, \succ)$; furthermore i is top ranked by o_i in the reduced problem. By overusing the notation, let $c =$

$\{o_k, i_k\}_{k=1,\ldots,K}$ be the cycle in which i is matched with $TTC(P, \succ)(i)$ and $i = i_K$. Furthermore o_{k+1} is i_k's most preferred object in the reduced problem. Repeating the arguments for c in the proof of Claim 1, we prove that φ and TTC assignments are the same for all agents in $\tilde{I}_2(P, \succ)$. Repeating this argument, we prove that φ and TTC assignments are the same for all agents in $\tilde{I}_k(P, \succ)$, all k.