

Pair-efficient reallocation of indivisible objects

ÖZGÜN EKICI

Department of Economics, Ozyegin University

We revisit the classical object reallocation problem under strict preferences. When attention is constrained to the set of Pareto-efficient rules, it is known that top trading cycles (TTC) is the only rule that is strategy-proof and individually rational. We relax this constraint and consider pair efficiency. A rule is pair-efficient if it never induces an allocation at which a pair of agents gain from trading their assigned objects. Remarkably, even in the larger set of pair-efficient rules, we find that TTC is still the only rule that is strategy-proof and individually rational. Our characterization result gives strong support to the use of TTC in object reallocation problems.

KEYWORDS. Indivisible object, pair efficient, strategy-proof, individually rational, top trading cycles.

JEL CLASSIFICATION. C78, D61, D63, D82.

1. INTRODUCTION

This paper considers the object reallocation problem: There is a group of agents, each of whom initially owns a distinct indivisible object. Agents have strict preferences over objects. Each agent's preference information is private. A rule specifies how to reallocate objects based on the preference information reported by agents.

We study object reallocation as a mechanism design problem: We are interested in rules satisfying desirable properties (axioms). We take two properties to be indispensable. First, the rule should incentivize agents to report preference information truthfully. We are interested in *strategy-proof* rules under which an agent never gains from misreporting. When this property is not satisfied, agents may strategize, which requires acquiring information about other agents' preferences and formulating better strategies. This process is arguably tiresome and wasteful. Second, we demand that the rule never assigns an agent an object worse than her endowment (the object that she owns). A rule satisfying this property is said to be *individually rational*. If a rule is not individually rational, agents may opt out. Therefore, individual rationality is a minimal voluntary participation constraint.

In an elemental result, Ma (1994) showed that when attention is constrained to the set of Pareto-efficient rules, the only strategy-proof and individually-rational rule is top

Özgün Ekici: ozgun.ekici@ozyegin.edu.tr

This paper greatly benefited from various comments and suggestions by Jay Sethuraman, William Thomson, and an anonymous referee. An abstract of the manuscript appeared in the proceedings of the 23rd ACM Conference on Economics and Computation (EC'22), July 11–15, 2022, Boulder, CO, USA.

trading cycles (TTC). Pareto efficiency is a natural efficiency requirement and, therefore, the result by Ma (1994) gives strong support to the use of TTC in object reallocation problems. This rule reallocates objects to agents in a stepwise manner by identifying and executing “top trading cycles.” A *top trading cycle*, or for short, a *cycle*, involves a group of agents who own one another’s favorite objects. When these cycles are executed, each agent involved in a cycle is assigned her favorite object. Then these agents and their assigned objects are removed from consideration, which leaves a reduced problem with a smaller number of agents and their endowments. In the reduced problem, new cycles are identified and executed similarly, and so on.

The result by Ma (1994) shows that to find strategy-proof and individually-rational rules other than TTC, one must relax the efficiency requirement. In this study, we do just that and substitute Pareto efficiency with pair efficiency. Pareto efficiency requires that at any induced allocation, no group of agents can gain from trading their assignments. In contrast, *pair efficiency* only requires that no pair of agents can gain from trading their assignments. Put differently, at an outcome, Pareto-efficiency rules out *every* efficiency-improving trade, but pair efficiency only rules out efficiency-improving trades involving *pairs* of agents. Since any efficiency-improving trade must involve at least two agents, pair efficiency is arguably a minimal efficiency requirement. We illustrate the extent to which pair efficiency relaxes Pareto efficiency in Example 1 in Section 2.2. In the example, for $n \geq 7$ agents, we describe a preference profile according to which there is a single Pareto-efficient allocation but the number of pair-efficient allocations exceeds 2^n .

Our relaxation of the Pareto-efficiency requirement also has a practical motivation. In this line of research, the literature focuses only on the welfare of the agents who trade objects, but we may imagine situations in which the social planner also has a stake in the outcome. For instance, when an employer (social planner) reallocates tasks (objects) to employees (agents), he may be interested in an outcome at which tasks will be performed productively. If the social planner cares about the outcome, the actual Pareto-efficient set of allocations becomes a superset of the set of allocations that are Pareto-efficient based only on agents’ preferences.

Indeed, for purposes of implementation, pair efficiency may be a more suitable requirement than Pareto efficiency. If an allocation is not Pareto-efficient, it admits an efficiency-improving trade. Therefore, after its implementation, agents could trade their assignments and destabilize the allocation. However, it would be a premature conclusion to say that a Pareto-inefficient allocation will *always* be destabilized. If the associated efficiency-improving trade cycles are all large, involving many agents, agents may find it hard to recognize and coordinate such trades. Therefore, after its implementation, the allocation may remain stable. However, the same argument cannot be made if an efficiency-improving trade involves only a pair of agents. A mutually beneficial pairwise exchange is easier to recognize and execute for the involved agents. Therefore, for implementation purposes, while Pareto efficiency may be too demanding, pair efficiency is a natural minimal requirement.

Although pair efficiency is a significant relaxation of Pareto efficiency, remarkably, the main result of our paper shows that in object reallocation, this relaxation does not give rise to new allocation rules. In Theorem 1, we show that TTC is characterized by

the properties of strategy-proofness, individual rationality, and pair efficiency. Put differently, even if the Pareto-efficiency requirement is relaxed and substituted by pair efficiency, TTC still turns out to be the unique strategy-proof and individually rational rule satisfying this property. As mentioned above, individual rationality is a minimal voluntary participation constraint and pair efficiency is a minimal efficiency requirement. Therefore, by showing that TTC is the only strategy-proof rule that satisfies these two minimal conditions, our main result gives very strong support for its use in object reallocation problems. The characterization result by Ma (1994) follows as a corollary of our Theorem 1.

The notion of pair efficiency has been explored in previous research in the context of the reallocation of divisible resources. Feldman (1973) explored the dynamics of pairwise barter trade to achieve a pairwise optimal allocation. Goldman and Starr (1982) introduced the more general t -wise optimality notion and developed necessary conditions and sufficient conditions for t -wise optimality to imply Pareto optimality. However, to our knowledge, ours is the first study that explores the concept in the context of the reallocation of indivisible resources.

We believe that a key contribution of our paper is our novel proof technique. Exploiting the procedural nature of TTC, we define an index that measures the level of similarity of outcomes induced by an arbitrary rule and TTC. This similarity index lies at the heart of our proof by minimal counterexample when showing Theorem 1. We believe that in future studies, working with a similarity index can also be useful while studying the properties of other procedural rules.

The rest of the paper is organized as follows. The following subsection presents other related research. Section 2 introduces the model and TTC. Section 3 presents our main result.

1.1 Other related research

The object reallocation problem was introduced by Shapley and Scarf (1974). The TTC rule is also first mentioned in their paper. They attributed it to David Gale and mentioned that it finds a core allocation. Roth and Postlewaite (1977) later showed that there is only one allocation in the core, which is found by TTC. Roth (1982) proved that TTC is strategy-proof; Bird (1984) showed that it is coalitionally strategy-proof.

There are several characterization studies on TTC in the literature. As mentioned above, Ma (1994) showed that TTC is the only rule that is strategy-proof, individually rational, and Pareto-efficient. Svensson (1994), Anno (2015), and Sethuraman (2016) provided shorter proofs of this result. Miyagawa (2002) showed that a rule that is strategy-proof, individually rational, anonymous, and nonbossy is either TTC or the endowment rule. In a more recent study, Fujinaka and Wakayama (2018) characterized TTC in terms of strategy-proofness, individual rationality, and endowments-swapping-proofness. This last property means that a pair of agents cannot both gain from trading their endowed objects *before* the rule is implemented. Notice that while pair efficiency is an efficiency criterion, endowments-swapping-proofness is a nonmanipulability notion. For two other characterization studies on TTC, see Takamiya (2001) and Morrill (2013).

In their paper, Hylland and Zeckhauser (1979) considered the object allocation problem, in which agents have no private endowments and objects are initially collectively owned. Abdulkadiroglu and Sönmez (1999) introduced the mixed-ownership extension of the object allocation problem. There is a line of research in the literature identifying classes of rules in object allocation and reallocation problems. In these studies, Pareto efficiency plays a pivotal role as the efficiency criterion: In a general class of allocation problems, Sönmez (1999) showed that there exists a Pareto-efficient, individually-rational, and strategy-proof solution only if the core is essentially single-valued (as in the object reallocation problem). In the mixed-ownership extension of the object allocation problem, Sönmez and Ünver (2010) characterized a class of rules by a set of properties that includes Pareto efficiency. In the object allocation problem, Pápai (2000) introduced hierarchical exchange rules and characterized them by the properties of Pareto efficiency, group-strategy-proofness, and reallocation-proofness. Later, Pycia and Ünver (2017) introduced an even larger class of rules called top cycles, and they characterized them by the properties of Pareto efficiency and group-strategy-proofness.

2. MODEL

2.1 Preliminaries

Let $I = \{1, 2, \dots, n\}$ be a finite set of agents. Let $O = \{o_1, o_2, \dots, o_n\}$ be a finite set of indivisible objects such that o_i denotes agent i 's endowment.

Agents are equipped with strict preferences over objects. Let $P = (P_i)_{i \in I}$ denote a preference profile where P_i denotes agent i 's strict preference relation. If agent i prefers object o to $\bar{o}$ at P_i , we write $o P_i \bar{o}$. Let R_i denote the *at least as good as* relation associated with P_i . Thus, $o R_i \bar{o}$ means $o P_i \bar{o}$ or $o = \bar{o}$.

When convenient, we describe a preference relation as an ordering of objects, from agent's first choice to last choice.

Let $\mathcal{P}$ be the set of strict preference relations over O . Thus, $P_i \in \mathcal{P}$ and $P \in \mathcal{P}^n$.

Sometimes, we work with mixed profiles. For $S \subseteq I$, $(\bar{P}_S, P_{I \setminus S})$ denotes the mixed profile such that for $i \in S$, the preference relation is $\bar{P}_i$ (as under profile $\bar{P}$), and for $i \in I \setminus S$, the preference relation is P_i (as under profile P). If $S = \{i\}$, we simply write $(\bar{P}_i, P_{I \setminus i})$. Later in the text, we work with a mixed profile $(P_{i_1}^+, P_{i_2}^\uparrow, P_{i_3}^\uparrow, \dots, P_{i_k}^\uparrow, P_{I \setminus S})$. This means that the preference relation is $P_{i_1}^+$ for i_1 , $P_{i_s}^\uparrow$ for $i_s \in \{i_2, \dots, i_k\}$, and P_i for $i \in I \setminus S$. Other mixed profile notations are understood accordingly.

An allocation assigns an object to each agent. A rule recommends an allocation for each preference profile. The formal definitions are as follows.

An *allocation* $\mu : I \rightarrow O$ is a one-to-one mapping from the set of agents to the set of objects. For $i \in I$, $\mu(i)$ denotes agent i 's assignment at μ . Let $\mathcal{M}$ denote the set of allocations.

A *rule* $\phi : \mathcal{P}^n \rightarrow \mathcal{M}$ is a mapping from the set of preference profiles to the set of allocations. That is, a rule ϕ associates each profile P with an allocation $\phi(P)$. For $i \in I$, $\phi_i(P)$ denotes agent i 's assignment at $\phi(P)$.

2.2 Axioms

We introduce next the axioms (properties) that are central to our analysis.

As a nonmanipulability condition, we consider strategy-proofness. Agents cannot manipulate (gain by misreporting) under a strategy-proof rule. The formal definition is as follows.

An agent i can manipulate a rule ϕ at profile P by misreporting her preferences as $\bar{P}_i \in \mathcal{P}$ ($\bar{P}_i \neq P_i$) if $\phi_i(\bar{P}_i, P_{I \setminus i}) P_i \phi_i(P)$. A rule ϕ is *strategy-proof* if no agent can manipulate it at any preference profile. That is, under a strategy-proof rule, reporting true preferences is a weakly dominant strategy.

As a voluntary participation condition, we consider individual rationality, which requires that an agent is never assigned an object worse than her endowment. Agents may opt out of an allocation rule if this condition is not satisfied. The formal definition is as follows.

An allocation μ is *individually rational* at P if for each $i \in I$, $\mu(i) R_i o_i$. A rule ϕ is *individually rational* if for each $P \in \mathcal{P}^n$, the allocation $\phi(P)$ is individually rational at P .

We will consider two efficiency notions. The first one is Pareto efficiency, which rules out at the outcome efficiency-improving trades between any group of agents. The second one is the weaker pair-efficiency notion, which only rules out efficiency-improving trades between pairs of agents. The formal definitions are as follows.

An allocation μ is *Pareto-efficient* at P if there exists no allocation $\bar{\mu}$ such that for each $i \in I$, $\bar{\mu}(i) R_i \mu(i)$, and for some $i \in I$, $\bar{\mu}(i) P_i \mu(i)$. A rule ϕ is *Pareto-efficient* if for each $P \in \mathcal{P}^n$, the allocation $\phi(P)$ is Pareto-efficient at P .

An allocation μ is *pair-efficient* at P if there do not exist $i, j \in I$, $i \neq j$, such that $\mu(i) P_j \mu(j)$ and $\mu(j) P_i \mu(i)$. A rule ϕ is *pair-efficient* if for each $P \in \mathcal{P}^n$, the allocation $\phi(P)$ is pair-efficient at P .

By definition, pair efficiency is a weaker notion than Pareto efficiency. We illustrate the extent to which pair efficiency relaxes Pareto efficiency in Example 1. In the example, we introduce a “circular” preference profile under which there is a single Pareto-efficient allocation, but the number of pair-efficient allocations exceeds 2^n when there are $n \geq 7$ agents.

EXAMPLE 1. We will consider a *circular* preference profile P^n under which objects can be labeled as $x_1, x_2, \dots, x_n$ such that for each agent i , $x_i P_i^n x_{i+1} P_i^n \dots P_i^n x_{i+n-1}$, where $x_{n+s} = x_s$. This is illustrated in Figure 1. Under P^n , objects can be placed around a circle




FIGURE 1. Preferences under a circular profile.

such that for each agent i , her preference ranking runs clockwise from her first choice x_i to her last choice x_{i-1} .

Under P^n , there is one Pareto-efficient allocation, where each agent i receives her first choice, x_i . Under P^n , let $F(n)$ be the number of pair-efficient allocations. We will show that $F(n) > 2^n$ for $n \geq 7$.

Let $f(n, s)$ be the number of pair-efficient allocations that assign exactly s out of n agents to their first choices. Then $F(n) = \sum_{s=0}^n f(n, s)$.

Next, we will calculate $f(n, s)$. Our calculation is recursive and depends on a key observation. Notice that under P^n , we can construct the pair-efficient allocations for which s agents receive their first choices in two steps as follows.

- We select s out of n agents and match them with their first choices.
- If $s < n$, we are left with a reduced problem with $n - s$ agents and $n - s$ objects. In the reduced problem, we choose a pair-efficient matching at which no agent receives her first choice. This matching, combined with the matches of s agents to their first choices, induces a pair-efficient allocation in our original problem with n agents and n objects.

Note that s out of n agents can be selected in $\binom{n}{s}$ ways. Also, in a problem with $n - s$ agents and $n - s$ objects, by definition, there are $f(n - s, 0)$ pair-efficient matchings at which no agent receives her first choice. Thus, $f(n, s) = \binom{n}{s} f(n - s, 0)$. To make this formula work for $s = n$, we set $f(0, 0) = 1$. Also, notice that, by definition, $f(1, 0) = f(2, 0) = 0$.

Thus, we get

$$F(n) = \sum_{s=0}^n \binom{n}{s} f(n - s, 0).$$

Let $n \geq 7$. We will also assume that n is odd. Thus, $n = 2k + 1$ for some $k \geq 3$. The assumption that n is odd is not essential for our arguments, but it helps simplify the exposition below. The interested reader can show that with a minor adjustment, our subsequent arguments can be used to show the desired result for n even, too.

Using the facts that $\binom{n}{n-s} = \binom{n}{s}$ and $\sum_{s=0}^n \binom{n}{s} = 2^n$, with some algebraic manipulation, we get

$$\begin{aligned} F(n) &= \sum_{s=0}^n \binom{n}{s} + \sum_{s=0}^n \binom{n}{s} (f(n - s, 0) - 1) \\ &= 2^n + \sum_{s=0}^k \binom{n}{s} (f(n - s, 0) + f(s, 0) - 2). \end{aligned}$$

To show that $F(n) > 2^n$, it suffices to show that in the above summation, the term $f(n - s, 0) + f(s, 0) - 2$ is always nonnegative and it is positive for $s = 0$.

One can easily verify that under a circular profile with $2t - 1$ or $2t$ agents where $t \geq 2$, we get a pair-efficient allocation when each agent receives her l th-best object for $l \in$

$\{2, 3, \dots, t\}$. Since $n \geq 7$, this implies the following inequalities: For $s = 0$, $f(n - s, 0) \geq 3$; for $s \in \{1, 2\}$, $f(n - s, 0) \geq 2$; for $s \in \{3, 4, \dots, k\}$, $f(n - s, 0) \geq 1$ and $f(s, 0) \geq 1$. Thus, as required, the term $f(n - s, 0) + f(s, 0) - 2$ is always nonnegative and it is positive for $s = 0$. Thus, $F(n) > 2^n$.

2.3 Top trading cycles

We introduce next the top trading cycles (TTC) rule.

A *top trading cycle*, or for short, a *cycle*, is a sequence of distinct agents and objects such that each agent in the cycle points to her favorite object and each object points to its owner. A cycle can be illustrated as

$$i_1 \rightarrow o_{i_2} \rightarrow i_2 \rightarrow o_{i_3} \rightarrow \dots \rightarrow i_{k-1} \rightarrow o_{i_k} \rightarrow i_k \rightarrow o_{i_{k+1}} \rightarrow i_{k+1} = i_1.$$

Above, $i_1, i_2, \dots, i_k$ are distinct agents. Agent i_s points to object o_{s+1} and o_{s+1} points to its owner i_{s+1} for $s = 1, \dots, k$. Thereby, the cycle forms.

The *size of a cycle* is the number of agents involved in that cycle. For instance, the size of the cycle indicated above is k . We write $|C|$ to indicate the size of a cycle C .

For two cycles C^1 and C^2 , we say that C^1 is *smaller* than C^2 if $|C^1| < |C^2|$ or if $|C^1| = |C^2|$ and in these two cycles, the agent whose index is smallest is in cycle C^1 . In a group of cycles, the *smallest* cycle is the one that is smaller than the others.

Note that given two cycles, if no agent is part of both cycles, one of them must be smaller than the other. Additionally, given a group of cycles, if no agent is part of multiple cycles, one cycle in the group must be the smallest. In the rest of the paper, when we make size comparisons, no agent will be part of multiple cycles. Therefore, in the groups of cycles that we consider, the smallest cycle will always be well defined. As an illustration, suppose that we are considering the group of cycles C^1, C^2, C^3 such that C^1 comprises the agents 2, 7, C^2 comprises the agents 3, 6, and C^3 comprises the agents 1, 4, 5. Then the group's smallest cycle is C^1 : C^1 is smaller than C^3 because $|C^1| = 2$ and $|C^3| = 3$; C^1 is also smaller than C^2 because $|C^1| = |C^2| = 2$, but in these two cycles, 2 is the agent whose index is smallest and it is part of C^1 .

When a cycle is *executed*, it means every agent involved in that cycle is assigned the object to which she points. For instance, if the cycle illustrated above is executed, agent i_s is assigned object $o_{i_{s+1}}$ for $s = 1, \dots, k$. When we say that an *allocation executes a cycle*, it means that at that allocation, every agent involved in that cycle is assigned the object to which she points. For instance, if allocation μ executes the cycle illustrated above, then $\mu(i_s) = o_{i_{s+1}}$ for $s = 1, \dots, k$.

The TTC rule, which is subsequently presented in a formal format, proceeds in a stepwise manner as follows. Every agent points to her favorite object and every object points to its owner. This gives rise to one or more cycles. These cycles are executed. In the next step, these agents and their assigned objects are removed from consideration. This leaves a reduced problem with a smaller number of agents and objects. The rule then operates on the reduced problem by identifying and executing new cycles, and so on.

It turns out that under TTC, the order in which cycles are executed is inconsequential.¹ For instance, let C^1 and C^2 be the cycles that arise at Step 1. One possibility is that we execute both C^1 and C^2 , and then proceed to Step 2. Alternatively, we can execute only C^1 and then proceed to Step 2. In this latter scenario, at Step 2, C^2 still remains. However, the execution of C^1 at Step 1 may trigger the formation of some new cycles, say C^3 and C^4 . Then we get three cycles at Step 2: C^2 , C^3 , C^4 . At Step 2, we may execute any combination of these three cycles and then proceed to Step 3. Ultimately, the same allocation is induced by TTC, independent of the order in which the cycles that arise are executed.

However, to show our main result, we need a specification of the TTC rule that exactly pinpoints which cycle is executed at which step. To this end, we will assume that at any point in time, TTC proceeds by executing only the smallest cycle (as defined above). We introduce this specification of TTC below.

Top Trading Cycles

Given a preference profile, TTC finds an allocation in a stepwise manner as follows.

Step 1. Construct a directed graph as follows. The nodes are agents and objects. For each agent, there is an edge from that agent to her favorite object. For each object, there is an edge from the object to its owner. Since there is an outgoing edge from each node and the nodes are finite, this gives rise to one or more cycles. Execute the smallest cycle among them. If a node remains, proceed to Step 2. Otherwise, terminate.

Step $t \geq 2$. With remaining agents and objects, construct a new directed graph as follows. The nodes are agents and objects. For each agent, there is an edge from that agent to her favorite (remaining) object. For each object, there is an edge from the object to its owner. Since there is an outgoing edge from each node and the nodes are finite, this gives rise to one or more cycles. Execute the smallest cycle among them. If a node remains, proceed to Step $t + 1$. Otherwise, terminate.

In the rest of the paper, when we speak of TTC, it is understood that we mean its above specification.

3. MAIN RESULT

THEOREM 1. *TTC is the only rule that is strategy-proof, individually rational, and pair-efficient.*

Ma (1994) showed that in the set of Pareto-efficient rules, TTC is the only rule that is strategy-proof and individually rational. Since pair efficiency is a relaxation of Pareto efficiency, his characterization result follows as a corollary of Theorem 1. In his paper, Ma (1994) also showed that strategy-proofness, individually rationality, and Pareto efficiency are independent axioms. That is, via three examples, he showed that a rule that satisfies two of these axioms need not satisfy the third one. The three examples in his paper can also be used to show the independence of the three axioms in Theorem 1. For the examples, the interested reader may refer to his paper.

The rest of the paper is devoted to the proof of Theorem 1. Section 3.1 describes our proof technique and introduces some tools. Section 3.2 presents our proof.

¹ See Remark 1 in Abdulkadiroglu and Sönmez (1999) and Lemma 6 in Carroll (2014).

3.1 Proof technique

To characterize a rule φ by a set of axioms, one must show that (i) φ satisfies these axioms and (ii) for an arbitrary rule ϕ , if ϕ satisfies the axioms, then $\phi = \varphi$. Most often, (ii) is the tricky part, which requires showing that $\phi(P) = \varphi(P)$ for every profile P . To show this, the *preference replacement technique* is often used, which proceeds as follows.

- For a specific profile P^1 , using the fact that ϕ satisfies the axioms, it is shown that $\phi(P^1) = \varphi(P^1)$.
- For an arbitrary profile P , a sequence of profiles $P^1, P^2, \dots, P^r$ is defined such that $P^r = P$. Using the fact that ϕ satisfies the axioms and given that $\phi(P^s) = \varphi(P^s)$, it is shown that $\phi(P^{s+1}) = \varphi(P^{s+1})$ for $s \in \{1, \dots, r-1\}$.

Our proof technique is different. To show the desired result, we introduce a TTC-similarity index function ρ . It measures the level of similarity of outcomes recommended by an arbitrary rule and TTC at a given preference profile. We write $\rho(\phi, P)$ for the TTC-similarity level of the outcomes $\phi(P)$ and $\text{TTC}(P)$. By definition, $\rho(\phi, P)$ becomes infinite if the two outcomes are the same and finite if otherwise. Then we show our main result by employing the proof by minimal counterexample method: If a rule ϕ that satisfies the axioms is not the same as TTC, $\rho(\phi, P)$ takes its smallest value for some profile P . We then modify P and construct a profile P^+ such that $\rho(\phi, P^+) < \rho(\phi, P)$, yielding a contradiction and proving the desired result.

The TTC-similarity index function ρ is defined by exploiting the TTC's procedural nature as follows. Consider running TTC with profile P .

- At Step 1, let C^1 be the cycle that is executed. If $\phi(P)$ does not execute C^1 , then $\rho(\phi, P) = (1, |C^1|)$. If $\phi(P)$ executes C^1 , check whether TTC terminates at Step 1. If yes, then $\rho(\phi, P) = (\infty, \infty)$. If not, proceed to Step 2.
- At Step t , let C^t be the cycle that is executed. If $\phi(P)$ does not execute C^t , then $\rho(\phi, P) = (t, |C^t|)$. If $\phi(P)$ executes C^t , check whether TTC terminates at Step t . If yes, then $\rho(\phi, P) = (\infty, \infty)$. If not, proceed to Step $t+1$.

Note that $\rho(\phi, P) = (t, k) \neq (\infty, \infty)$ means the following. At profile P , ϕ assigns objects first by running TTC with profile P and executing the smallest cycles that arise at Steps 1, 2, $\dots$, $t-1$. However, ϕ then deviates from TTC. Specifically, $\phi(P)$ does not execute the smallest cycle that arises at Step t , which is of size k .

We compare the TTC-similarity levels using the lexicographic order as follows. For $P, \bar{P} \in \mathcal{P}^n$, let $\rho(\phi, P) = (x_1, y_1)$ and $\rho(\phi, \bar{P}) = (x_2, y_2)$. We write $\rho(\phi, P) \leq \rho(\phi, \bar{P})$ if $x_1 < x_2$, or if $x_1 = x_2$ and $y_1 \leq y_2$. We write $\rho(\phi, P) < \rho(\phi, \bar{P})$ if $\rho(\phi, P) \leq \rho(\phi, \bar{P})$ and $\rho(\phi, P) \neq \rho(\phi, \bar{P})$.

We also define a TTC-similarity level for a rule ϕ without reference to a preference profile. We denote it by $\rho(\phi)$ and set its value equal to the minimum value that $\rho(\phi, P)$ takes on the set of preference profiles. That is, $\rho(\phi) = \rho(\phi, P)$ where $P \in \mathcal{P}^n$ is such that $\rho(\phi, P) \leq \rho(\phi, \bar{P})$ for all $\bar{P} \in \mathcal{P}^n$.

Note that if $\rho(\phi) = (t, k) \neq (\infty, \infty)$, it means that for any profile P , ϕ makes assignments first by running TTC with profile P and executing the smallest cycles that arise at

Steps $1, 2, \dots, t-1$, but then ϕ deviates from TTC for at least one profile. Specifically, there is a profile P such that when TTC runs with P , the smallest cycle at Step t is of size k and $\phi(P)$ does not execute this cycle.

Also, note that if $\rho(\phi) = (\infty, \infty)$, it means that for any profile P , $\phi(P) = \text{TTC}(P)$. That is, to show that an arbitrary rule ϕ is the same as TTC, we need to show that $\rho(\phi) = (\infty, \infty)$.

We are now ready to present our proof.

3.2 Proof of Theorem 1

PROOF. It is known that TTC is strategy-proof, individually rational, and pair efficient. Thus, we will only show that if a rule ϕ satisfies these three axioms, then $\phi = \text{TTC}$.

Suppose by contradiction that ϕ satisfies these three axioms but $\phi \neq \text{TTC}$. Then $\rho(\phi) \neq (\infty, \infty)$. Let $\rho(\phi) = (t, k)$ for finite integers t, k .

Let $P \in \mathcal{P}^n$ be such that $\rho(\phi, P) = (t, k)$. Consider TTC running with profile P . Let $C^1, C^2, \dots, C^t$ be, in order, the sets of cycles that arise at Steps $1, 2, \dots, t$. Let $C^1 \in \mathcal{C}^1$, $C^2 \in \mathcal{C}^2, \dots, C^t \in \mathcal{C}^t$ be, in order, the cycles that are executed at Steps $1, 2, \dots, t$.

Since $\rho(\phi, P) = (t, k)$, $\phi(P)$ executes the cycles $C^1, C^2, \dots, C^{t-1}$, but not C^t . Let cycle C^t be

$$i_1 \rightarrow o_{i_2} \rightarrow i_2 \rightarrow o_{i_3} \rightarrow \dots \rightarrow i_{k-1} \rightarrow o_{i_k} \rightarrow i_k \rightarrow o_{i_{k+1}} \rightarrow i_{k+1} = i_1.$$

Let $S = \{i_1, i_2, \dots, i_k\}$. We use i_s to denote an arbitrary agent in S . Note that when objects involved in cycles $C^1, C^2, \dots, C^{t-1}$ are excluded, among remaining objects, o_{s+1} is the favorite object at P_{i_s} .

We will prove the theorem by deriving a contradiction to our supposition that $\rho(\phi) = (t, k)$. We will do this by constructing a preference profile P^+ for which $\rho(\phi, P^+) = (t, k-1)$. We will construct P^+ from P by changing the preference relations of agents in S one at a time. Along the process, we will consider the execution of TTC with these new profiles. Specifically, we will compare the cycles that arise and that are executed at Steps $1, 2, \dots, t$ when TTC runs with these profiles versus when it runs with profile P . In our arguments, two observations pertaining to these cycles will become useful. We will present them next.

First, notice that when TTC runs with profile P , up to and including Step t , an agent $i_s \in S$ never points to an object that she ranks below o_{s+1} . Therefore, the execution of TTC up to this point will not change if i_s changes her preference ordering of objects below o_{s+1} . We state this observation formally below.

Observation 1. Consider a profile $(\bar{P}_S, P_{I \setminus S})$ such that for each $i_s \in S$, objects ordered above o_{s+1} are the same, and are ordered in the same way at $\bar{P}_{i_s}$ and P_{i_s} . Then, when TTC runs with profile $(\bar{P}_S, P_{I \setminus S})$, in order, $C^1, C^2, \dots, C^t$ are the sets of cycles that arise and $C^1, C^2, \dots, C^t$ are the cycles that are executed, at Steps $1, 2, \dots, t$.

For our next observation, consider an arbitrary profile P^+ . When TTC runs with P^+ , suppose that the cycles that arise at Steps $1, 2, \dots, t$ are the same cycles that arise at these steps when TTC runs with profile P , but with the following exceptions. When

TTC runs with P , we know that the cycle C^t arises at some step and then is executed at Step t . However, the cycle C^t never arises when TTC runs with profile P^+ . Instead, a new cycle, C^+ , arises at some step, where $|C^+| < |C^t|$, and C^+ remains unexecuted until Step t . However, if it is so, notice that running TTC with P^+ instead of P will not lead to a change in the cycles that are executed at Steps $1, 2, \dots, t-1$. Removing C^t from and adding C^+ to the previous cycles is inconsequential since neither C^t nor C^+ becomes a smallest cycle at Steps $1, 2, \dots, t-1$. Also, notice that when TTC runs with P^+ , C^+ becomes the executed cycle at Step t since $|C^+| < |C^t|$ and C^t is the smallest cycle in C^t . We state this observation formally below.

Observation 2. When TTC runs with profile P^+ , let $\bar{C}^1, \bar{C}^2, \dots, \bar{C}^t$ be, in order, the sets of cycles that arise at Steps $1, 2, \dots, t$ such that for $s = 1, 2, \dots, t-1$, $\bar{C}^s \setminus \{C^+\} = C^s \setminus \{C^t\}$ and $\bar{C}^t = C^t \setminus \{C^t\} \cup \{C^+\}$. Then, when TTC runs with profile P^+ , in order, $C^1, C^2, \dots, C^{t-1}, C^+$ are the cycles that are executed at Steps $1, 2, \dots, t$.

Having made these two observations, we are now ready to proceed with the rest of our proof.

Since $\phi(P)$ does not execute C^t , at least one agent in this cycle is not assigned the object to which she points. Without loss of generality, let i_k not be assigned o_1 at $\phi(P)$. Thus, $\phi(P)$ assigns i_k an object that is worse than o_{i_1} at P_{i_k} . By individual rationality of ϕ , $k \neq 1$. Thus, $k \geq 2$.

In the remainder of the proof, we will consider certain preference relations. We illustrate them below.

P_{i_s}	$P_{i_s}^\uparrow$	$P_{i_1}^*$	$P_{i_1}^+$
$\vdots$	$\vdots$	$\vdots$	$\vdots$
$o_{i_{s+1}}$	$o_{i_{s+1}}$	o_{i_2}	o_{i_3}
$\vdots$	o_{i_s}	o_{i_3}	o_{i_1}
o_{i_s}	$\vdots$	o_{i_1}	o_{i_2}
$\vdots$	$o_{i_{s+2}}$	$\vdots$	$\vdots$
	$\vdots$		

- Above, $P_{i_s}^\uparrow$ is obtained from P_{i_s} by moving o_{i_s} up, right below $o_{i_{s+1}}$. The order for $P_{i_s}^\uparrow$ includes in it $o_{i_{s+2}}$ as if $k \geq 3$. However, disregard $o_{i_{s+2}}$ if $k = 2$.
- Above, $P_{i_1}^*$ and $P_{i_1}^+$ are defined assuming that $k \geq 3$: $P_{i_1}^*$ is obtained from $P_{i_1}^\uparrow$ by moving o_{i_3} up, right below o_{i_2} ; $P_{i_1}^+$ is obtained from $P_{i_1}^*$ by moving o_{i_2} down, right below o_{i_1} .

Consider the profile $P^k = (P_{i_k}^\uparrow, P_{I \setminus i_k})$. If $\phi_{i_k}(P^k)P_{i_k}^\uparrow o_{i_k}$, then $\phi_{i_k}(P^k)R_{i_k} o_{i_{k+1}}$. Recall that $\phi(P)$ assigns i_k an object worse than $o_{i_{k+1}}$ at P . However, then i_k can manipulate ϕ at P by reporting $P_{i_k}^\uparrow$. This cannot be true since ϕ is strategy-proof. Then, also using the fact that ϕ is individually rational, we get $\phi_{i_k}(P^k) = o_{i_k}$. By Observation 1, when TTC runs with profiles P^k and P , the same sets of cycles arise, and the same cycles are executed, at Steps $1, 2, \dots, t$. Since $\rho(\phi) = (t, k)$, $\phi(P^k)$ executes the cycles $C^1, C^2,$

$\dots, C^{t-1}$. When objects assigned in the cycles $C^1, C^2, \dots, C^{t-1}$ are excluded, among remaining objects, i_{k-1} 's favorite object at $P_{i_{k-1}}$ is o_{i_k} . Since $\phi_{i_k}(P^k) = o_{i_k}$, $\phi(P^k)$ assigns i_{k-1} an object worse than o_{i_k} at $P_{i_{k-1}}$.

Consider the profile $P^{k-1} = (P_{i_{k-1}}^\uparrow, P_{i_k}^\uparrow, P_{I \setminus \{i_{k-1}, i_k\}})$. If $\phi_{i_{k-1}}(P^{k-1})P_{i_{k-1}}^\uparrow o_{i_{k-1}}$, then $\phi_{i_{k-1}}(P^{k-1})R_{i_{k-1}} o_{i_k}$. However, then i_{k-1} can manipulate ϕ at P^k by reporting $P_{i_{k-1}}^\uparrow$. This cannot be true since ϕ is strategy-proof. Then, also using the fact that ϕ is individually rational, we get $\phi_{i_{k-1}}(P^{k-1}) = o_{i_{k-1}}$. By Observation 1, when TTC runs with profiles P^{k-1} and P , the same sets of cycles arise, and the same cycles are executed, at Steps 1, 2, $\dots, t$. Since $\rho(\phi) = (t, k)$, $\phi(P^{k-1})$ executes the cycles $C^1, C^2, \dots, C^{t-1}$. Suppose $k = 2$. However, then, by individual rationality of ϕ , we get $\phi_{i_k}(P^{k-1}) = o_{i_k}$. Then, at $\phi(P^{k-1})$, according to reported preferences, i_k and i_{k-1} prefer one another's assigned objects to their own assignments. This cannot be true since ϕ is pair-efficient. Thus, $k \geq 3$. When objects assigned in the cycles $C^1, C^2, \dots, C^{t-1}$ are excluded, among remaining objects, i_{k-2} 's favorite object at $P_{i_{k-2}}$ is $o_{i_{k-1}}$. Since $\phi_{i_{k-1}}(P^{k-1}) = o_{i_{k-1}}$, $\phi(P^{k-1})$ assigns i_{k-2} an object worse than $o_{i_{k-1}}$ at $P_{i_{k-2}}$.

As described above, we can iteratively define the profiles $P^{k-2}, P^{k-3}, \dots, P^1$. By applying similar arguments, for profile $P^1 = (P_{i_1}^\uparrow, P_{i_2}^\uparrow, \dots, P_{i_k}^\uparrow, P_{I \setminus S})$, we find the following: $\phi_{i_1}(P^1) = o_{i_1}$, and when TTC runs with profiles P^1 and P , the same sets of cycles arise, and the same cycles are executed at Steps 1, 2, $\dots, t$. However, then, by individual rationality of ϕ , we find that for each $i_s \in S$, $\phi_{i_s}(P^1) = o_{i_s}$.

Consider now the profile $P^* = (P_{i_1}^*, P_{i_2}^\uparrow, P_{i_3}^\uparrow, \dots, P_{i_k}^\uparrow, P_{I \setminus S})$. By Observation 1, when TTC runs with profiles P^* and P , the same sets of cycles arise and the same cycles are executed at Steps 1, 2, $\dots, t$. Since $\rho(\phi) = (t, k)$, $\phi(P^*)$ executes the cycles $C^1, C^2, \dots, C^{t-1}$. When objects assigned in the cycles $C^1, C^2, \dots, C^{t-1}$ are excluded, among remaining objects, i_1 's favorite three objects are, in order, $o_{i_2}, o_{i_3}, o_{i_1}$. By individual rationality of ϕ , we get $\phi_{i_1}(P^*) \in \{o_{i_2}, o_{i_3}, o_{i_1}\}$.

If $\phi_{i_1}(P^*) = o_{i_2}$, then i_1 can manipulate ϕ at P^1 by reporting $P_{i_1}^*$. This cannot be true since ϕ is strategy-proof. Thus, $\phi_{i_1}(P^*) \neq o_{i_2}$.

If $\phi_{i_1}(P^*) = o_{i_3}$, by individual rationality of ϕ , we get $\phi_{i_2}(P^*) = o_{i_2}$. But then, at $\phi(P^*)$, according to reported preferences, i_1 and i_2 prefer one another's assigned objects to their own assignments. This cannot be true since ϕ is pair-efficient. Thus, $\phi_{i_1}(P^*) \neq o_{i_3}$.

Thus, $\phi_{i_1}(P^*) = o_{i_1}$.

Consider now the profile $P^+ = (P_{i_1}^+, P_{i_2}^\uparrow, P_{i_3}^\uparrow, \dots, P_{i_k}^\uparrow, P_{I \setminus S})$. Since ϕ is strategy-proof, we get $\phi_{i_1}(P^+) = o_{i_1}$.

By definition of P^+ , when TTC runs with profile P^+ , the cycle C^t does not arise. Instead, at Step t or perhaps earlier, the cycle C^+ ,

$$i_1 \rightarrow o_{i_3} \rightarrow i_3 \rightarrow o_{i_4} \rightarrow \dots \rightarrow i_{k-1} \rightarrow o_{i_k} \rightarrow i_k \rightarrow o_{i_{k+1}} \rightarrow i_{k+1} = i_1,$$

arises.

Since $\phi_{i_1}(P^+) = o_{i_1}$, $\phi(P^+)$ does not execute the cycle C^+ . However, then, since $\rho(\phi) = (t, k)$, when TTC runs with profile P^+ , the cycle C^+ remains unexecuted until Step t .

We found the following occurrences. When TTC runs with profile P , the cycle C^t arises at some step and it remains unexecuted until Step t , when it is executed. When TTC runs with profile P^+ , the cycle C^+ arises at some step and it remains unexecuted until Step t . However, notice that when the cycles C^t and C^+ are put aside, the remaining cycles that arise at Steps $1, 2, \dots, t$ are the same when TTC runs with profiles P and P^+ . Therefore, Observation 2 is applicable and, hence, when TTC runs with profile P^+ , C^+ is executed at Step t . Since $\phi(P^+)$ does not execute C^+ , we get $\rho(\phi, P^+) = (t, k - 1)$, which contradicts that $\rho(\phi) = (t, k)$. This proves that $\phi = \text{TTC}$. $\square$

REFERENCES

- Abdulkadiroglu, Atila and Tayfun Sönmez (1999), “House allocation with existing tenants.” *Journal of Economic Theory*, 88, 233–260. [554, 558]
- Anno, Hidekazu (2015), “A short proof for the characterization of the core in housing markets.” *Economics Letters*, 66–67. [553]
- Bird, Charles G. (1984), “Group incentive compatibility in a market with indivisible goods.” *Economics Letters*. [553]
- Carroll, Gabriel (2014), “A general equivalence theorem for allocation of indivisible objects.” *Journal of Mathematical Economics*, 51, 163–177. [558]
- Feldman, Allan M. (1973), “Bilateral trading processes, pairwise optimality, and Pareto optimality.” *The Review of Economic Studies*, 40, 463–473. [553]
- Fujinaka, Yuji and Takuma Wakayama (2018), “Endowments-swapping-proof house allocation.” *Games and Economic Behavior*, 111, 187–202. [553]
- Goldman, Steven M. and Ross M. Starr (1982), “Pairwise, t-wise, and Pareto optimalities.” *Econometrica*, 50, 593. [553]
- Hylland, Aanund and Richard Zeckhauser (1979), “The efficient allocation of individuals to positions.” *Journal of Political Economy*, 87, 293–314. [554]
- Ma, Jinpeng (1994), “Strategy-proofness and the strict core in a market with indivisibilities.” *International Journal of Game Theory*, 23, 75–83. [551, 552, 553, 558]
- Miyagawa, Eiichi (2002), “Strategy-proofness and the core in house allocation problems.” *Games and Economic Behavior*, 38, 347–361. [553]
- Morrill, Thayer (2013), “An alternative characterization of top trading cycles.” *Economic Theory*, 54, 181–197. [553]
- Papai, Szilvia (2000), “Strategyproof assignment by hierarchical exchange.” *Econometrica*, 68, 1403–1433. [554]
- Pycia, Marek and M. Utku Ünver (2017), “Incentive compatible allocation and exchange of discrete resources.” *Theoretical Economics*, 12, 287–329. [554]
- Roth, Alvin E. (1982), “Incentive compatibility in a market with indivisible goods.” *Economics Letters*, 9, 127–132. [553]

Roth, Alvin E. and Andrew Postlewaite (1977), “Weak versus strong domination in a market with indivisible goods.” *Journal of Mathematical Economics*, 4, 131–137. [553]

Sethuraman, Jay (2016), “An alternative proof of a characterization of the TTC mechanism.” *Operations Research Letters*, 44, 107–108. [553]

Shapley, Lloyd and Herbert Scarf (1974), “On cores and indivisibility.” *Journal of Mathematical Economics*, 1, 23–37. [553]

Sönmez, Tayfun and M. Utku Ünver (2010), “House allocation with existing tenants: A characterization.” *Games and Economic Behavior*, 69, 425–445. [554]

Svensson, Lars Gunnar (1994), “Queue allocation of indivisible goods.” *Social Choice and Welfare*. [553]

Takamiya, Koji (2001), “Coalition strategy-proofness and monotonicity in Shapley–Scarf housing markets.” *Mathematical Social Sciences*, 41, 201–213. [553]

Co-editor Federico Echenique handled this manuscript.

Manuscript received 16 October, 2022; final version accepted 5 July, 2023; available online 17 July, 2023.