Economics Bulletin

Volume 46, Issue 2

An impossibility in housing markets with contracts

Jinyong Jeong
Li Anmin Institute of Economic Research, Liaoning University

Abstract

This paper extends the classical housing-market model of Shapley and Scarf (1974) to an environment with contractual terms and demonstrates an impossibility result: no mechanism can simultaneously satisfy Pareto efficiency, term-consistent property rights, and strategy-proofness. The result isolates how introducing contractual dimensions disrupts the alignment of incentives that underlies the top trading cycles mechanism. The analysis identifies the minimal structure generating nonexistence and clarifies the limits of truthful implementation in exchange economies with indivisible goods and contract-contingent rights.

Citation: Jinyong Jeong, (2026) "An impossibility in housing markets with contracts", *Economics Bulletin*, Volume 46, Issue 2, pages 767-778
Contact: Jinyong Jeong - jinyong.t.jeong@gmail.com.
Submitted: October 07, 2025. **Published:** June 30, 2026.

1. Introduction

The Shapley and Scarf (1974) housing market and the top trading cycles (TTC) algorithm deliver a unique core outcome that is Pareto efficient, individually rational, and strategy-proof in the classical setting (Ma, 1994; Roth, 1982; Roth and Postlewaite, 1977). Recent work has further refined our understanding of TTC. For example, Klaus, Klijn, and Sethuraman (2025) provide a new characterization via respecting-improvement.

In many applications, however, objects are exchanged under explicit *contractual terms* (e.g., duration, restrictions), and feasible reallocations must respect these terms. Building on the matching-with-contracts perspective (Hatfield and Milgrom, 2005), we consider a housing market in which each house can be traded under two terms and where an agent's endowment imposes a term-consistency constraint. For a comprehensive overview of these developments, Abdulkadiroğlu and Sönmez (2013) provide a well-structured and informative review of this line of research. This integration also draws inspiration from the concept of matching with property rights outlined in Jeong (2021). While the focus of Jeong (2021) is primarily on stability, our paper places a key emphasis on efficiency as one of its solution criteria.

The framework of this paper is related to but conceptually distinct from Moulin (1995)'s multi-dimensional housing model. While Moulin (1995) considers houses with multiple inherent attributes (location, size, type), with the impossibility arising from heterogeneous preferences over independent dimensions. In contrast, we focus on contractual terms that impose contingent restrictions on property use and transfer. The equal-term constraint encodes a property-rights consistency principle that is absent in Moulin's framework: the term under which an agent receives a house must match the term under which her endowment is used.

Our main result contributes to a growing literature on impossibility results in matching markets. Recent studies have identified fundamental trade-offs between efficiency, fairness, and incentive compatibility in various settings (Ramezanian and Feizi (2022)). In the context of mechanism design with constraints, Imamura and Kawase (2025) analyze general constraint structures. Our impossibility result shows that no mechanism can achieve (i) Pareto efficiency, (ii) an equal-term property that encodes term-consistent property rights, and (iii) strategy-proofness. We then discuss how positive procedures—such as dual variants of TTC that operate term-by-term—recover efficiency and equal-term but inevitably fail strategy-proofness. The takeaway is that introducing contracts breaks the alignment of desiderata enjoyed in the classical TTC environment.

2. Model

There are n agents $I = \{1, \dots, n\}$ and n houses $H = \{h_1, \dots, h_n\}$, with agent i initially endowed with h_i. A *contract* is a triple (i, h, t), where $i \in I$, $h \in H$, and $t \in \{R, L\}$ denotes a contractual term. Each agent i has strict preferences $\succ_i$ over acceptable contracts and the outside option $\bar{i} = (i, h_i, R)$ (keeping her house under the default R term). A *preference profile* is a list $\succ = (\succ_1, \dots, \succ_n)$. We write $\succ_{-i}$ for the profile of all agents other than i.

A *matching* μ assigns each agent a single contract, with no house assigned to more than one agent. Formally, $\mu : I \to I \times H \times \{R, L\}$ with $\mu(i) = (i, h, t)$ for some $h \in H$ and $t \in \{R, L\}$, and $\mu(i) \neq \mu(j)$ whenever $i \neq j$.

Definition 1 (Feasibility and IR) *A matching μ is* feasible *if each house $h \in H$ appears in at most one assigned contract. It is* individually rational *(IR) if $\mu(i) \succeq_i \bar{i}$ for all $i \in I$.*

Definition 2 (Equal-term property) [1] *A matching μ satisfies* equal-term *if, for every agent i, the term of $\mu(i)$ equals the term of the contract involving house h_i. That is, letting $t(\mu(i))$ denote the term component of $\mu(i)$ and $\mu^{-1}(h_i)$ the agent assigned h_i, we require $t(\mu(i)) = t(\mu(\mu^{-1}(h_i)))$ for all i. Equivalently, the term attached to an agent's assignment equals the term under which her endowment is used.*

Definition 3 (Efficiency and strategy-proofness) *A feasible μ is* (Pareto) efficient *if there is no feasible ν with $\nu(i) \succeq_i \mu(i)$ for all i and $\nu(j) \succ_j \mu(j)$ for some j. A* mechanism *φ maps each preference profile $\succ$ to a feasible matching $\varphi(\succ)$. It is* strategy-proof *if for every agent i, every profile $\succ$, and every misreport $\succ_i'$, we have $\varphi(\succ)(i) \succeq_i \varphi(\succ_i', \succ_{-i})(i)$.*

We study mechanisms that map preference profiles to feasible matchings.

3. Main Impossibility

We consider three desiderata for a mechanism's outcome: (E) efficiency, (ET) equal-term, and (SP) strategy-proofness. The following shows they are mutually incompatible.

Proposition 1 *No mechanism can simultaneously satisfy (E), (ET), and (SP) in housing markets with contracts.*

Proof: Consider agents 1 and 2 with strict preferences

$$1: \ (2, L) \succ_1 (2, R) \succ_1 \bar{1}, \qquad 2: \ (1, R) \succ_2 (1, L) \succ_2 \bar{2}.$$

We first characterize all feasible equal-term matchings. Since equal-term requires the term of agent i's assignment to match the term under which h_i is used, in a two-agent market where the agents trade houses, both must receive contracts of the same term. There are thus exactly two feasible IR equal-term matchings:

$$\mu^L = \big((2, L), (1, L)\big) \quad \text{and} \quad \mu^R = \big((2, R), (1, R)\big),$$

where the i-th coordinate is the contract assigned to agent i. (The outside options $\bar{1}$ and $\bar{2}$ are also equal-term, but they are Pareto dominated by both μ^L and μ^R since every agent strictly prefers trading her house to keeping it.)

[1] In a general context, being **equal-term** implies **upholding property rights** defined in (Jeong, 2021), and the two are equivalent when each agent possesses their own house.

Let φ be any mechanism satisfying (E) and (ET). Then $\varphi(\succ) \in \{\mu^L, \mu^R\}$ for the profile $\succ$ above.

Case 1: $\varphi(\succ) = \mu^L$. Agent 2 receives $(1, L)$ but strictly prefers $(1, R)$. Suppose agent 2 misreports $\succ_2'$ by declaring $(1, L)$ unacceptable (i.e., ranking only $(1, R)$ above $\bar{2}$). Under $(\succ_1, \succ_2')$, the matching μ^L is no longer IR for agent 2, so any mechanism satisfying (E) and (ET) must select μ^R. Agent 2 strictly prefers $\mu^R(2) = (1, R) \succ_2 (1, L) = \mu^L(2)$, so φ is not strategy-proof.

Case 2: $\varphi(\succ) = \mu^R$. Agent 1 receives $(2, R)$ but strictly prefers $(2, L)$. Suppose agent 1 misreports $\succ_1'$ by declaring $(2, R)$ unacceptable (ranking only $(2, L)$ above $\bar{1}$). Under $(\succ_1', \succ_2)$, the matching μ^R is no longer IR for agent 1, so any mechanism satisfying (E) and (ET) must select μ^L. Agent 1 strictly prefers $\mu^L(1) = (2, L) \succ_1 (2, R) = \mu^R(1)$, so φ is not strategy-proof.

Since $\varphi(\succ)$ must be either μ^L or μ^R, in both cases (SP) fails. ■

Remark 1 *The construction is minimal: the conflict is generated by two agents, two houses, and two contractual terms, with only the equal-term constraint added relative to the classical Shapley–Scarf model. Since a mechanism must produce an efficient outcome for every preference profile, this single counterexample suffices to establish the impossibility. Dropping the equal-term constraint restores compatibility via TTC.*

4. Discussion

Role of equal-term. Equal-term encodes a simple property-rights consistency: the term under which an agent receives a house must match the term under which her endowment is used. Proposition 1 shows that this seemingly innocuous constraint is sufficient to break the TTC property. The key insight is that equal-term creates a coupling between an agent's assignment and the use of her own endowment, which generates a conflict between the agent's incentives and social efficiency. When two agents disagree on the preferred term—one preferring L and the other preferring R—any efficient equal-term mechanism must choose one term, and the agent who receives the less-preferred term can always profitably misreport to shift the outcome.

Positive mechanisms under weaker goals. One can recover efficiency and equal-term via procedures that operate *term-by-term* (e.g., dual variants of TTC that clear equal-term cycles, detailed in the appendix). However, such procedures necessarily sacrifice strategy-proofness: agents can manipulate which term-clearing path is triggered by misreporting their preference ranking over terms.[2]

To illustrate the nature of strategy-proofness failure more concretely, consider a market where agent i prefers an L-term contract but knows that truthfully reporting this will trigger an L-term clearing cycle that yields her a less preferred house. By misreporting a lower

[2]See the appendix for the detailed mechanism.

preference for L-term contracts, she may instead trigger an R-term cycle that assigns her a more preferred house under R-term, which she may still prefer overall to the L-term outcome. This type of cross-term manipulation is unavoidable whenever agents have preferences over both the house and the contractual term.

Conversely, in restricted domains where all agents strictly prefer one common term, the TTC mechanism applied to that single term is both efficient and strategy-proof, but then the equal-term constraint is vacuously satisfied, and the model reduces to the classical Shapley–Scarf setting

Efficiency relaxation and partial recovery. An alternative response to the impossibility is to relax efficiency rather than strategy-proofness. One such relaxation is *constrained efficiency*: a matching is constrained efficient if there is no other equal-term matching that Pareto dominates it. A strategy-proof and constrained-efficient mechanism exists whenever one restricts attention to equal-term matchings from the outset and applies TTC within this restricted domain. However, constrained efficiency may impose welfare losses relative to full efficiency: agents may fail to realize mutually beneficial trades that would require mixing terms. The welfare cost of this restriction depends on the distribution of preferences over terms and houses—in markets where agents are nearly homogeneous in their term preferences, the cost is small, while in markets with strong disagreement over terms, the loss can be substantial.

Core selection. In classical housing markets, TTC is a core mechanism. With contracts, the equal-term core may be empty even in small markets, and when non-empty, any core-selecting rule inherits manipulability in the sense of Proposition 1.

5. Conclusion

Introducing contractual terms into housing markets creates a fundamental conflict among efficiency, equal-term consistency, and strategy-proofness. The impossibility holds for any number of agents and identifies the precise cost of enforcing term-consistent property rights: either truthful implementation must be abandoned, or efficiency must be relaxed. The dual-TTC mechanism in the appendix shows that efficiency and equal-term consistency can be jointly achieved, at the cost of strategy-proofness. Future work can map domains where partial recoveries are possible (e.g., limited preference interdependencies across terms) and quantify the welfare cost of strategy-proof relaxations.

References

Abdulkadiroğlu, Atila and Tayfun Sönmez (2013). "Matching markets: theory and practice." *Advances in Economics and Econometrics*. Vol. 1. Cambridge University Press, pp. 3–47.

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

Imamura, Kenzo and Yasushi Kawase (2025). "Efficient and strategy-proof mechanism under general constraints." *Theoretical Economics*, 20 (2), 481–509.

Jeong, Jinyong (2021). "Matching with property rights: an application to korean teacher transfer program." *Review of Economic Design*, 25, 139–156.

Klaus, Bettina, Flip Klijn, and Jay Sethuraman (2025). "A characterization of the top-trading-cycles mechanism for housing markets via respecting-improvement." *Economics Letters*, 247, 112145.

Ma, Jinpeng (1994). "Strategy-proofness and the strict core in a market with indivisibilities." *International Journal of Game Theory*, 23, 75–83.

Moulin, Hervé (1995). *Cooperative Microeconomics: A Game-Theoretic Introduction*. Princeton University Press.

Ramezanian, Rasoul and Mehdi Feizi (2022). "Robust ex-post Pareto efficiency and fairness in random assignments: Two impossibility results." *Games and Economic Behavior*, 135, 356–367.

Roth, Alvin E. (1982). "Incentive compatibility in a market with indivisible goods." *Economics Letters*, 9, 127–132.

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

Shapley, Lloyd and Herbert Scarf (1974). "On cores and indivisibility." *Journal of Mathematical Economics*, 1 (1), 23–37.

A. Dual Top Trading Cycles

To find an efficient equal-term matching, I suggest the implementation of a dual-TTC algorithm. This algorithm aims to identify equal-term matchings in both R-term and L-term contracts.

Before introducing the algorithm, I define several terms used in describing the process. Throughout, we fix a market $(I, H, \succ)$ with $|I| = |H| = n$.

Definition 4 *A **loop** is an ordered list of agents $(i_1, i_2, \cdots, i_k)$ such that agent i_j demands agent i_{j+1}'s house for each $j = 1, \ldots, k-1$, and agent i_k demands agent i_1's house.*

*An **equal-term-loop** is a loop $(i_1, \ldots, i_k)$ in which each agent demands the house of the next agent under contracts of the same term $t \in \{R, L\}$; that is, for each j, the demanded contract is $(i_j, h_{i_{j+1 \pmod k}}, t)$ for a common t.*

*A **top rank contained equal-term-loop** is an equal-term-loop in which at least one agent's demand within the loop is her most preferred available contract.*

In contrast to the original housing market, the dual-TTC algorithm may involve multiple overlapping cycles, making it essential to determine how to select one among them. Cycle choice rules will be assumed to be exogenously given as they will depend on the specific application of the model. Without specifying the cycle choice rule, I define the new algorithm as follows.

Definition 5 *Dual-Top Trading Cycles*[3]

step 1 : Each agent points to her most preferred contract among all acceptable L-term contracts and her most preferred contract among all acceptable R-term contracts. If an agent exclusively has one type of acceptable contracts, she points to only one contract of that type. Proceed to step 2.

*step 2 : If there is any **top rank contained equal-term-loop**, clear one loop based on the choice rule by assigning each agent to their requested contract and removing all the agents in the loop, along with all the contracts associated with these agents. Return to step 1. Otherwise, proceed to step 3.*

step 3 : If there is at least one agent whose acceptable contracts are of one type, proceed to step 3-1. Otherwise, proceed to step 3-2.

step 3-1 : Temporarily remove the agent who is not part of the loop and has only a single term of available contracts. Then, the agent who was initially pointing to the removed agent now points to the next available contract. Continue this process until a top rank contained equal-term-loop forms, and return to step 2. If no such loop is created even after all agents are removed in this phase, proceed to step 3-2.

[3] In the appendix, readers can find the algorithm illustrated using figures.

step 3-2 : All the loops in this step are not top rank contained. Randomly pick one agent and remove her top rank contract. Note that the removed contract does not form a loop. Then, the agent now points to the next acceptable contract of the same term. Continue with another agent until a top rank contained loop forms. If no such loop is created even after all agents removed their choices that are not part of the loop, then clear the equal-term-loop without a top rank. Return to step 1.

The algorithm terminates when there are no acceptable contracts for every agent, or there is no loop that forms with remaining acceptable contracts.

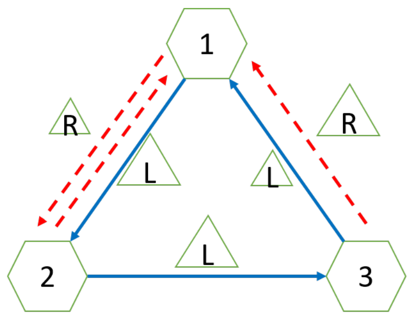

Figure 1: Dual-Top Trading Cycles

To establish that the dual TTC is a well-defined algorithm and is both efficient and individually rational, the following lemmas are useful.

Lemma 1 *If every agent possesses at least one of each R-term and L-term acceptable contract, then there is at least one equal-term-loop.*

Proof: When every agent has both R-type and L-type acceptable contracts, we can split the market into two sub-markets: one for R-type contracts and the other for L-type contracts. In each sub-market, every agent points to exactly one contract (her top-ranked contract of that type). Since the sub-market has finitely many agents, the resulting directed graph—in which each agent points to exactly one other agent (the holder of her demanded house)—must contain a cycle. Since all contracts within each sub-market share the same term, each such cycle is an equal-term-loop. ■

Lemma 2 *If there is no top rank contained equal-term-loop, then temporarily removing an agent who has only a single type of acceptable contracts and is not part of any loop does not lead to an efficiency loss in the final matching.*

Proof: See Figure 2 for an illustrative example. Let i^* be an agent with only R-term acceptable contracts, and suppose there is no top rank contained equal-term-loop. We argue that i^*'s top-ranked contract cannot be part of any top rank contained equal-term-loop.

Since i^* has only R-term contracts, she points only within the R-term sub-market. For i^*'s demand to be part of a top rank contained equal-term-loop, there must exist an R-term loop that includes i^*. But if such a loop existed and i^*'s demand is her top rank, then this loop would be a top rank contained equal-term-loop, contradicting the assumption. Hence i^*'s top-ranked demand is not achievable and she cannot be part of any improving equal-term trade.

Now temporarily remove i^* from the market. By Lemma 1, the remaining agents (who by assumption all have both R and L acceptable contracts) admit at least one equal-term-loop. Clearing a top rank contained loop among the remaining agents does not affect i^*, since i^* cannot be part of any such loop. Moreover, the agents in the cleared loop receive their demanded contracts, which are efficient for them given that at least one agent in the loop is receiving her top-ranked contract. Hence no efficiency loss occurs from temporarily removing i^*. ■

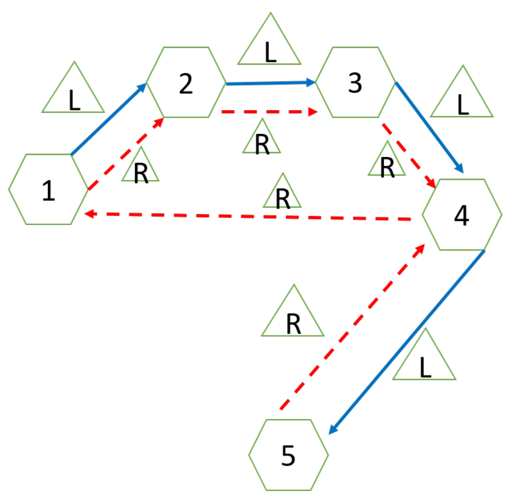

Figure 2: No top rank contained equal-term-loop. Agent i_5 has only R-type contract in this case. If i_5 had an acceptable L-type contract, there should have been a top rank contained equal-term-loop.

In rare cases, there may not be a top rank contained equal-term-loop, even when everyone has both R and L acceptable contracts, as shown in the following example:

Table 1: Table I: Preference profile with no top rank contained equal-term-loop

i_1	i_2	i_3	i_4	i_5
$(4, L)$	$(4, L)$	$(4, L)$	$(1, R)$	$(2, R)$
$(2, R)$	$(3, R)$	$(1, R)$	$(5, L)$	$(4, L)$
$\bar{1}$	$\bar{2}$	$\bar{3}$	$\bar{4}$	$\bar{5}$

There are two equal-term-loops: one among i_1, i_2, and i_3 with R-term contracts, and the other between i_4 and i_5 with L-term contracts, and neither is top rank contained. Lemma 3 establishes that removing one top-ranked contract from an agent who has an alternative of the same type does not cause efficiency loss.

Lemma 3 *If there is no top rank contained equal-term-loop but every agent has both R and L-term acceptable contracts, then removing one of the top-ranked contracts from an agent who has another acceptable contract of the same term does not result in efficiency loss.*

Proof: Suppose there is no top rank contained equal-term-loop and every agent has both R and L-term acceptable contracts. Let i be any agent and suppose her top-ranked t-term contract (for $t \in \{R, L\}$) is (i, h_j, t) for some $j \neq i$. Since there is no top rank contained equal-term-loop, this contract is not part of any such loop, meaning that it cannot be simultaneously the top-ranked contract for agent i within a loop and be used in clearing an efficient trade.

We claim that removing (i, h_j, t) from i's acceptable set does not prevent any efficient equal-term assignment from being realized. To see this, suppose for contradiction that some efficient equal-term matching μ^* assigns contract (i, h_j, t) to agent i, and that μ^* cannot be realized after removing (i, h_j, t).

Since (i, h_j, t) is not part of any top rank contained equal-term-loop, it follows that when agent i demands (i, h_j, t), no loop forms in which all agents simultaneously obtain their top-ranked demands. Therefore, realizing μ^* requires some agents to accept contracts that are not their top-ranked choice. In particular, for μ^* to be efficient, it must be that no agent can improve upon μ^* within the equal-term feasible set.

Now consider the matching obtained by assigning agent i her second-ranked t-term contract instead. Since i's top-ranked t-term contract is unavailable in any top rank contained loop, the second-ranked t-term contract can be used to form an alternative loop structure. By the assumption that no top rank contained equal-term-loop exists, any equal-term loop that clears must involve agents who are not at their top-ranked contract anyway. Removing (i, h_j, t) simply redirects i to her next preference, and the loop structure that eventually forms clears the same set of mutually beneficial trades. Hence no efficiency loss occurs. ■

Lemma 4 *If there is no equal-term-loop after removing agents with no acceptable contracts, then there must be at least one agent whose acceptable contracts are all of L-type, and at least one agent whose acceptable contracts are all of R-type. Moreover, in this case there is no equal-term loop that is individually rational for all agents within it.*

Proof: Suppose, after removing all agents with no acceptable contracts, there is no equal-term-loop. We show that the remaining agents cannot all have both R and L-term acceptable contracts.

By Lemma 1, if every remaining agent had both R and L-term acceptable contracts, there would exist at least one equal-term-loop. Since we assumed there is no such loop, it must be that some agents have only one type of acceptable contract.

We now show that there must be at least one agent with only L-type contracts and at least one with only R-type contracts. Suppose for contradiction that all remaining agents with a single type all have the same type, say R. Consider the directed graph G_R on I in which agent i points to the holder of her top-ranked R-term contract (if she has one). Since all agents either have both types or only R-type, G_R is defined for all remaining agents. The graph G_R has n nodes each with out-degree 1, so it must contain a directed cycle. This cycle is an equal-term R-loop, contradicting our assumption. Hence we cannot have all single-type agents of the same type; both types must be present.

Finally, we argue that no IR equal-term loop exists. An IR loop requires every agent in the loop to weakly prefer her assigned contract to her outside option $\bar{i}$. But with no equal-term loop at all (by assumption), there is trivially no IR equal-term loop. ■

i_1	i_2	i_3	i_4
$(2, L)$	$(3, R)$	$(4, L)$	$(2, R)$
$(3, L)$	$\boxed{(3, L)}$		$(1, R)$
$\boxed{(4, R)}$	$(1, L)$		
$\bar{1}$	$\bar{2}$	$\bar{3}$	$\bar{4}$

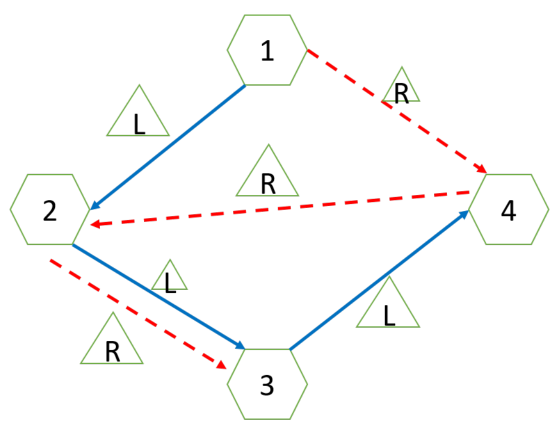

Figure 3: No equal-term-loop has formed. i_3 has only L-type acceptable contract, and i_4 has only R-type acceptable contract.

Table 2: Table II: Preference profile illustrating Lemma 4

i_1	i_2	i_3	i_4
$(2, L)$	$(3, R)$	$(4, L)$	$(2, R)$
$(3, L)$	$\boxed{(3, L)}$		$(1, R)$
$\boxed{(4, R)}$	$(1, L)$		
$\bar{1}$	$\bar{2}$	$\bar{3}$	$\bar{4}$

Theorem 1 *The dual TTC produces an efficient, individually rational, and equal-term matching.*

Proof: Dual TTC produces an individually rational and equal-term matching since all agents only demand contracts that are acceptable to them at each step, and the algorithm clears only equal-term-loops.

We prove that Dual TTC is efficient by showing that each step does not violate efficiency when it clears the forming loops. First, step 2 produces an efficient exchange since it contains at least one most preferred contract for an agent within the loop.

By Lemma 2, the first part of step 3-1 does not violate efficiency when removing an agent who has a single term of acceptable contracts. Then by Lemma 1, there should be at least one equal-term-loop. If this process forms a top rank contained equal-term-loop, then clearing such a loop is efficient. If there still does not exist a top rank contained one, then, by Lemma 3, we can proceed to step 3-2 without violating efficiency.

Lastly, by Lemma 4, step 3-2 is still not violating efficiency when it removes contracts from agents. Note that the algorithm returns to the previous steps depending on the resulting loops, which does not violate efficiency either. ■

Dual TTC is an efficient mechanism, but it is hardly strategy-proof. To achieve strategy-proofness, one would need rather extreme assumptions about the underlying preferences of agents, such as all agents sharing the same strict ordering over terms.