Top Trading Cycles in Prioritized Matching:
An Irrelevance of Priorities in Large Markets

Yeon-Koo Che* Olivier Tercieux†

August 27, 2018

Abstract

Top trading cycles (TTC) provides a method for assigning resources efficiently while taking agents' priorities into account. Although TTC favors agents with high priorities in the assignment, we show in a canonical random preference/priority model that the effect of priorities in TTC disappears as the market grows large, leading in the limit to an assignment that entails virtually the same amount of justified envy as does Random Serial Dictatorship, which completely ignores priorities.

JEL Classification Numbers: C70, D47, D61, D63.
Keywords: Prioritized matching, Markov property, irrelevance of priorities in TTC

1 Introduction

In prioritized resource allocation, agents are assigned indivisible objects based on their preferences as well as their priorities for objects. Prioritized allocation accounts for many important allocation problems, including assignment of students to public schools, of tenants to public housing, and of human organs to transplant patients. The priorities of agents typically take the form of priority ranks at each object,[1] reflecting their "merits" or other public policy consideration. In school assignment, for instance, a student's standardized test score (in the case of exam schools) or walk-zone or sibling attendance (in the case of regular schools) may

*Department of Economics, Columbia University, USA. Email: yeonkooche@gmail.com.
†Department of Economics, Paris School of Economics, France. Email: tercieux@pse.ens.fr.

[1]Therefore, formally, agents' priorities for an object can be treated equivalently as the object's preference ranking of the agents.

determine one's priority status. Pareto efficient allocation is a clear desideratum in any resource allocation problem, but it is also viewed as an important goal to respect participants' priorities, or equivalently to eliminate *justified envy*—a situation in which an agent envies another despite having a higher priority at the latter's assignment (Balinski and Sönmez (1999) and Abdulkadiroglu and Sonmez (2003)).[2]

If efficiency is the key objective,[3] there are two commonly used methods for assignment: Random Serial Dictatorship (RSD) and (prioritized) Top Trading Cycles (TTC) algorithm. RSD randomly orders participants and let them pick their most preferred remaining objects at their turns. TTC, introduced by Abdulkadiroglu and Sonmez (2003), has in each round agents point to their most preferred objects and objects point to their highest priority agents, forms "cycles" of agents and objects, and assigns the agents in the cycles the objects they point to.[4] Both methods achieve Pareto efficiency, have been used in practice,[5] and are strategy-proof; i.e., no agents have incentives to misreport their preferences.

The key distinction of TTC in comparison with RSD is its use of agents' priorities in the allocation. Unlike RSD which completely ignores agents' priorities, TTC explicitly favors agents with high priorities in allocation. Specifically, at any round, if an agent has the highest priority for his most preferred remaining object, then he is guaranteed to obtain that object under TTC. And, if he indeed obtains that object, no priorities for that object can be violated, or equivalently he cannot be justifiably envied by any other agents.[6] This feature makes one hopeful that TTC may do well in eliminating justified envy, at least among those that satisfy Pareto efficiency and strategyproofness. Indeed, Abdulkadiroglu, Che, Pathak, Roth, and Tercieux (2017) demonstrate that in one-to-one matching TTC is *justified envy minimal* within the mechanisms satisfying the two properties—namely no other efficient and strategyproof mechanism reduces the set of agents with justified envy.[7] They further show that for each

[2]Respecting priorities is also part of stability requirement of Gale and Shapley (1962) and amounts to elimination of justified envy.

[3]Efficiency is incompatible with respecting priorities (Roth (1982)). If the latter is the key objective, then a Gale and Shapley's agent-proposing deferred acceptance algorithm achieves that goal with the minimal sacrifice of efficiency; i.e., it outputs an non-wasteful assignment that respects priorities and Pareto dominates all other non-wasteful assignments respecting priorities.

[4]More detailed description will be given later. Note that the TTC we focus on differs from the well-known Shapley-Scarf TTC, where the primitives involve ownership of objects but no priorities (see Shapley and Scarf (1974)).

[5]RSD is used in school assignment and college dorm assignment. TTC was used until recently in New Orleans school systems for assigning students to public high schools. A generalized version of TTC is also used for kidney exchange among donor-patient pairs with incompatible donor kidneys (see Sonmez and Unver (2011)).

[6]Anybody who envies that agent must be remaining at that round of TTC, so must have a lower priority for that object.

[7]By contrast, the RSD is not justified envy minimal in the same sense.

profile of preferences, TTC admits fewer incidences of justified envy *on average* than RSD, when the average is taken over all possible profiles of priorities.[8]

Just "how well" does TTC do in respecting priorities? Namely what is the "quantitative" significance of the benefit TTC yields in respecting priorities say over and above RSD? The aforementioned characterization is silent on this question. A main contribution of this paper is to demonstrate that, at least in a canonical environment, the answer to this question is in the negative.

Consider one-to-one matching with n agents and n objects where the agents' preferences and their priorities for objects are drawn iid uniformly. We show that as the economy grows large with $n \to \infty$, the outcome of TTC in terms of the *joint* distribution of agents' preference ranks *and* their priorities becomes asymptotically equivalent to that under the RSD.[9] This means, among other things, that TTC does virtually no better in respecting priorities than does RSD. More precisely, the proportion of agents with justified envy (or those whose priorities are respected) under TTC becomes indistinguishable from that under RSD, both from the average and probabilistic senses when the market grows large.

The reason for this striking result is explained by the particular way in which TTC uses agents' priorities for allocating objects. If an object is assigned via a *short cycle*—or a cycle in which an agent points to an object and the object points back to that agent—, then it is in fact impossible that an agent's priority for that object is violated. The matters are quite different, however, if an object is assigned via a *long cycle*, or a cycle of length more than two. In that case, the acquiring agent has no *a priori* reason to have a high priority for that object; the only reason for the agent to have acquired that object is that he has a high priority for "some other" object he has traded off. Indeed, any envy for an agent assigned via a long cycle is likely to be "justified" with probability one half, just as with RSD. Our irrelevance follows from the fact that as $n \to \infty$, the proportion of objects assigned via short cycles vanishes in probability. Of course, in most applications, the distribution of preferences and priorities is not uniform iid. As we argue in Section 6, however, our result holds for fairly large classes of distributions.

In practice, RSD is used widely in a variety of contexts such as public housing assignment, college housing assignment, and many resource allocation problems without transfer. There is a potential policy interest in replacing RSD with TTC in such contexts, on the ground that

[8]Without taking an average, this comparison between TTC and RSD fails. For instance, one can find profiles of priorities and preferences for which TTC admits more (expected) incidences of justified envy than does RSD. See Abdulkadiroglu, Che, Pathak, Roth, and Tercieux (2017).

[9]The well-known equivalence result of Pathak and Sethuraman (2011) and Carroll (2014) (extending Abdulkadiroglu and Sönmez (1998)) means that the *marginal* distribution of the ranks enjoyed by agents are identical between the two mechanisms. As will be noted below, however, the *joint* distribution of agents' preference ranks *and* their priorities are distinct in finite economies.

TTC may protect individuals' priorities better than RSD. A striking implication of our result is that in a large market resembling our environment, one should not expect a switch to TTC to make a significant difference.

While our irrelevance result appears intuitive, its proof requires a deep and precise characterization of how the TTC allocates objects in our random model. A crucial step is to establish a Markov property: the number of objects assigned at any round of TTC follows a simple Markov chain, with the number depending only on the number of agents and objects at the beginning of that round in a well-specified manner. The Markov characterization allows us to show that TTC algorithm terminates in the number of rounds which is sublinear in n. With the expected number of objects assigned via short cycles further shown never to exceed two per round, this implies that the proportion of objects assigned via short cycles vanishes in probability, leading ultimately to the irrelevance claim stated above.

We view the Markov characterization of TTC as our second main contribution, of independent value beyond the particular application explored in the current paper. It is of interest since it can lead to a precise understanding of the outcome of TTC, in terms of the distribution of agents' preference ranks and their realized priority ranks. While the former is known from the analysis of RSD due to its equivalence,[10] the current analysis may shed additional light on the latter based on the Markov characterization.[11]

The current paper is related to several strands of literature. First, it is related to the literature studying the tradeoff between efficiency and elimination of justified envy particularly in the school choice context. The tradeoff was first recognized by Roth (1982), and was confirmed by Abdulkadiroglu and Sonmez (2003) in the school choice context, and by Abdulkadiroglu, Pathak, and Roth (2009) in the context of weak priorities. Che and Tercieux (2015) argue that the tradeoff remains significant under standard mechanisms (such as DA or TTC) even in a large market if there is a sufficient correlation in agents' preferences. Abdulkadiroglu, Che, Pathak, Roth, and Tercieux (2017) show the sense in which TTC minimizes justified envy/maximizes priority-respecting in the class of Pareto-efficient and strategyproof mechanisms in one-to-one matching setting. Together with the irrelevance result of the current paper, this result yields the sense in which the irrelevance is driven by the Pareto efficiency and strategyproofness rather than the feature of TTC itself.

Second, the irrelevance result is closely related to the equivalence among well-known random allocations recognized by a number of authors (Knuth (1996), Abdulkadiroglu and Sönmez (1998), Pathak and Sethuraman (2011), Carroll (2014), Bade (2016)). These authors

[10]See Knuth (1996) for the rank distribution of agents under RSD and see Pathak and Sethuraman (2011) for the equivalence which generalizes that of Knuth (1996) and Abdulkadiroglu and Sönmez (1998).

[11]For instance, the distribution of realized priority ranks has not been known before. We provide a characterization of the ranks that lead to the irrelevance result mentioned above, as well as the asymptotic instability of TTC in the large market studied in Che and Tercieux (2015).

show that a class of random allocations, including TTC and RSD, implement the identical (probabilistic) assignment of agents to objects. However, they are silent on the *joint* distribution of agents' preference ranks *and* their priorities under alternative mechanisms. As illustrated, the joint distribution of allocation matters for the extent to which agents' priorities are respected and the extent to which agents justifiably envy others. As we will shortly illustrate, the alternative mechanisms are not equivalent in this regard for a finite market. Nevertheless, the equivalence is restored as the market grows large. Our result therefore can be seen as the strengthening of equivalence (to include priority rank distribution), for the iid preferences case.

Third, the Markov characterization of TTC is related closely to a similar Markov characterization of the Shapley-Scarf TTC derived in Frieze and Pittel (1995).[12] In the Shapley-Scarf economy, the agents are endowed with *property rights* over objects, *exactly one object for each agent*, and are allowed to trade their rights along cycles in successive rounds. Despite the close resemblance, the two mechanisms are distinct. The associated ("pointing") map from objects to agents is always bijective in the Shapley-Scarf TTC but not in the prioritized TTC; distinct objects typically point to the same agent. This difference leads to different probabilistic structures in the associated composite map—agents pointing to objects which in turn point to agents—in our random economy, and required different arguments albeit following a similar approach. The concepts of random spanning forests and random composite maps prove crucial in our analysis, which to our knowledge have never been applied in economics. We believe they will constitute a useful tool in other economic applications.

The remainder of the paper is organized as follows. Section 2 illustrates the main irrelevance result in an example. Section 3 introduces the formal model and preliminary tools for analysis. Section 4 provides the Markov characterization. Section 5 presents the irrelevance result and its implications. Section 6 discusses robustness and limitation of our results; we suggest in particular that the irrelevance result holds much more generally beyond the uniform iid draws of preferences and priorities.

2 Example

Suppose there are two agents, 1 and 2, and two objects, a and b, where agents' (ordinal) preferences over objects and objects' (ordinal) priorities over agents are all drawn iid uniformly.

[12] Another related work is Leshno and Lo (2017), which studies TTC in a large market but with a very different asymptotics where the number of object types is finite while there are a continuum of copies/seats for each object type has and a continuum of agents with finite preference types. This distinction makes the analysis largely unrelated. Nevertheless, we discuss how our irrelevance result extends to this alternative asymptotics (see Section 6).

As noted earlier, both RSD and TTC allocate the agents efficiently. Further, the equivalence result of Pathak and Sethuraman (2010) and Carroll (2014) means that for each realized profile of preferences, the agents enjoy identical lotteries from both mechanisms. However, this equivalence does not mean that the joint distribution of allocation is identical between the two. In fact, their allocations differ in terms of the realized priorities and justified envy.

To see this, note first that TTC eliminates justified envy completely. When the agents prefer distinct objects, then Pareto efficiency simply means that they get their preferred objects, so no envy exists. Suppose therefore the agents prefer the same objects, say a. Under TTC, both agents point to a in round 1, and a points to the agent with a higher priority for a. This means that the higher priority agent obtains a. So envy toward that agent is never justified.

The same is not true for RSD. Whenever both agents prefer the same object, again say a, there is one half chance that the low priority agent gets the top serial order and claims a, in which case the high priority agent has justified envy.

Indeed, the expected value of realized priority ranks differs between the two mechanisms. Under RSD, the realized priority ranks are completely uniform random, so the expected value of realized priority rank equals $\text{RANK}_{RSD} = 3/2(= \frac{1}{2} \cdot 1 + \frac{1}{2} \cdot 2)$. By contrast, under TTC, the expected value of the realized priority ranks is lower, equalling $\text{RANK}_{TTC} = 11/8$.[13] The superior performance of TTC relative to RSD in eliminating justified envy is a consequence of Abdulkadiroglu, Che, Pathak, Roth, and Tercieux (2017), and hence holds generally, regardless of the market size.

Nevertheless, the *magnitude* of the superior performance by TTC over RSD—the focus of the current paper—diminishes as the market grows. Suppose that the number of agents and objects equal $n \geq 2$, and the agents' preferences and their priorities for objects are drawn iid uniformly. Panel (a) of Figure 1 plots the total incidences of justified envy among the incidences of envy under RSD and TTC, across different market sizes.[14] Panel (b) of Figure 1

[13] This can be seen as follows. When the agents prefer different objects, which occurs with probability 1/2, the agents obtained their preferred objects, so their priority ranks are completely (uniform) random just like RSD. When both agents prefer the same object (which occurs with probability 1/2), that object is assigned to the top-priority agent, and the other object is assigned uniform randomly in terms of the realized priority. In sum, an object has probability 1/4 of being the commonly preferred item, in which case it is assigned to a top priority agent, and has probability 3/4 of not being the commonly preferred item in which case the priority rank realized is the same as under RSD, i.e., 3/2 on average.

$$\text{RANK}_{TTC} = \tfrac{1}{4} \cdot 1 + \tfrac{3}{4} \cdot \tfrac{3}{2} = 11/8.$$

[14] More precisely, we sum across all agents the number of individuals that they justifiably envy and divide this by the total number of envies—not necessarily justifiable—experienced by the agents. Hence, the current measure reflects the "intensity" of justified envy experienced by the agents, instead of just counting the

plots the expected value of realized (normalized) priority ranks under the two mechanisms.[15]

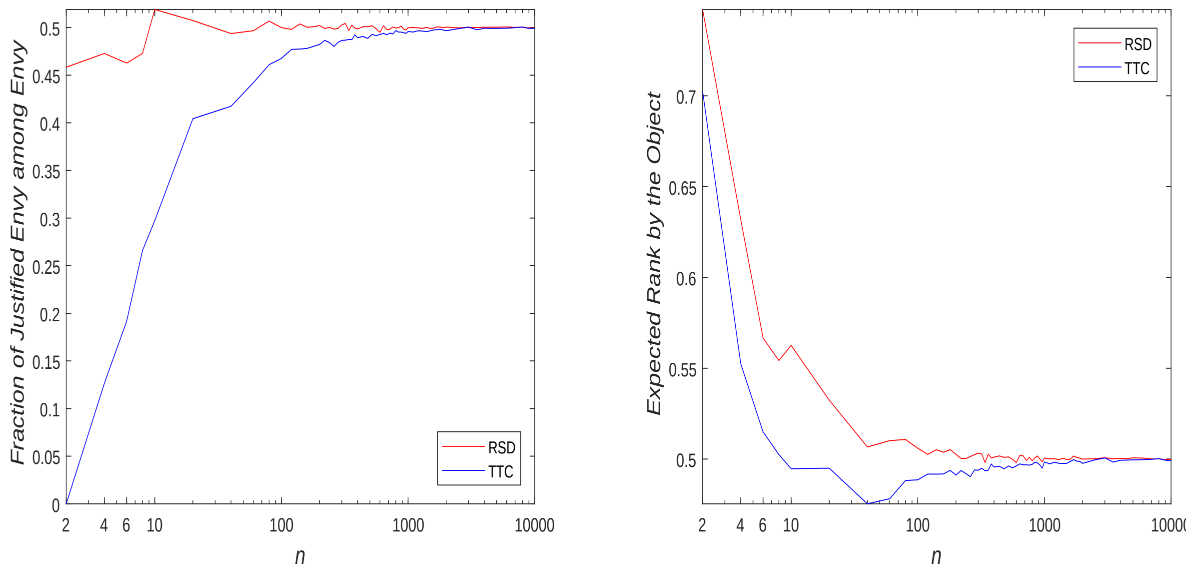

(a) Incidences of justified among incidences of envy

(b) Expected value of realized priority ranks

Figure 1: Difference between TTC and RSD.

Note: The horizontal axis indicates the size of the market, i.e., the number of agents/objects. The figures plot averages over 100 (uniform iid) draws of agents' preferences and objects' priorities.

In keeping with Abdulkadiroglu, Che, Pathak, Roth, and Tercieux (2017), both quantities are higher under RSD than under TTC. However, the differences between the two mechanisms diminishes as the market grows large. Ultimately, the difference between the two mechanisms is traced to the use of agents' priorities (or objects' ranking of agents) in organizing short cycles (i.e., cycles in which an agent points to an object and the object points directly back to him). The difference is large with $n = 2$ due to the importance of short cycles. Figure 1 illustrates that as the market grows large, the proportion of objects assigned via short cycles decreases. And this implies that the use of priorities under TTC becomes increasingly irrelevant as the market grows large.

proportion of agents with the justified envy. Of course, we could use this latter measure or the fraction of blocking pairs out of all pairs. While the results are similar in terms of these other measures, the current measure is more informative, since the fraction of blocking pairs converges to zero even for RSD (when the number of objects equals that of the agents) and the fraction of agents with justified envy converges to one even for TTC. The measure is also more favorable to TTC and thus stacks the deck against irrelevance, as will become clear in Section 6.

[15]That is, for each mechanism, we take the average (across objects) of the priority ranks enjoyed by the objects and divide it by the number of agents for normalization.

3 Model

We consider a market with a set I of **agents** and a set O of **objects**. We do not require $|I| = |O|$, so the market could be unbalanced. A **matching** is a map $\mu : I \to O \cup \{\emptyset\}$ such that $i \neq i'$ means that either $\mu(i) \neq \mu(i')$ or $\mu(i) = \mu(i') = \emptyset$, where $\emptyset$ means being unmatched.

In Section 5, we will consider a large market, and for that purpose, we will assume $|I| = |O| = n$ and consider the outcome as $n \to \infty$. The preference P_i of agent $i \in I$ is a strict ranking, or a permutation, of O. The agents' priorities for object $o \in O$, denoted by $\succ_o$, is a strict ranking, or a permutation, of I. As mentioned earlier, one can equivalently view agents' priorities for object o as its preference ranking of them, although this will not be part of welfare consideration. We assume that all entities on each side are acceptable to those on the other side.[16] Let $(P, \succ) := (P_i, \succ_o)_{i\in I, o\in O}$ be a profile of agents' preferences and their priorities. Throughout, we shall consider a **random market** $(I, O, \tilde{P}, \tilde{\succ})$ in which the profile $(\tilde{P}, \tilde{\succ})$ is drawn *iid* uniformly.

There are at least two ways to interpret random preferences and priorities in a real-world context such as school choice. One can think of a cohort of applicants for schools in each year as a result of random drawing of student types. What we are studying is then the properties of TTC that are likely to hold for such a cohort with a high probability as the market grows large. While the uniform iid assumption may not seem realistic given this interpretation, we explain in Section 6 how our results extend beyond the uniform iid case. Another interpretation is to view the randomness as an analytical device for taking an average. For instance, the random market approach will allow us to consider average of performance measures such as justified envy over all possible profiles of students' preferences and priorities.

In the case of RSD, there is an additional randomness due to the use of random lottery, which is a feature of the mechanism itself. To accommodate this, we define ω to be a **state** which consists of a profile $(P, \succ)$ of preferences on both sides as well as a random serial order θ, a permutation of I. A **mechanism** is then a mapping from each state to a matching. Given our random economy, randomness may arise from the random preferences/priorities or the random serial order θ, as is the case with RSD.

Fix a realized state $\omega = (P, \succ, \theta)$. RSD and TTC are defined as follows:

□ **RSD:** At step 1, the agent $\theta(1)$ at the top order of θ claims object at the top of his preference $P_{\theta(1)}$, and exits the market. At step $t \geq 2$, the agent $\theta(t)$ at the t-th serial order of θ claims the most preferred object among those remaining at t according to $P_{\theta(t)}$. This algorithm terminates, and all the agents are fully assigned, at step n. The RSD mechanism selects a matching via this algorithm for all possible profiles of agents' preferences and their

[16] This is for convenience. Our result holds more generally if the number of acceptable partners grow linearly in n.

priorities. Note that RSD only uses (P, θ) but ignores $\succ$.

$\square$ **TTC:** In Round $t = 1$ each individual $i \in I$ points to his most preferred object according to P_i. Each object $o \in O$ points to the individual who has the highest priority for that object according to $\succ_o$. Since the number of individuals and objects are finite, the directed graph thus obtained has at least one cycle. Every individual who belongs to a cycle is assigned the object he points to. All assigned individuals and objects are then removed. In round $t \geq 2$, the same procedure is taken among the agents and objects remaining at the beginning of that round. The algorithm terminates when all individuals or all objects have been assigned; otherwise, it proceeds to Round $t + 1$. This mechanism terminates in finite rounds, since there is a finite set of agents, and at least one individual is removed at the end of each round. The TTC mechanism selects a matching via this algorithm for all possible profiles of agents' preferences and their priorities. Note that TTC uses both P and $\succ$ (and ignores θ).

As noted, both RSD and TTC mechanisms are **Pareto efficient**; that is, for each profile P of agents' preferences, the matching produced by either mechanism cannot be improved upon by a different matching that makes all agents weakly better off and some strictly better off. Note that Pareto efficiency is defined only taking the agents' welfare into account. Next, both mechanisms are **strategyproof**, meaning that it is a dominant strategy for each agent to report truthfully.[17]

In fact, the two mechanisms are identical from the agents' perspectives. The two random matching mechanisms induced for each profile P of agents' preferences give rise to an identical lottery for each agent (Pathak and Sethuraman (2010); Carroll (2014)). As noted in Section 2, the equivalence does not extend to the allocation from the perspective of the objects. To obtain more precise comparison of the mechanisms in this regard, we need to understand the probabilistic structure of allocation in TTC.

4 Markov Chain Property of TTC

Our first main result is a Markov chain characterization of TTC: the numbers of agents and objects that are assigned in each round of TTC follow a simple Markov chain depending *only* on the numbers of agents and objects at the beginning of that round, according to well-defined transition probabilities. The theorem is stated in a general setting with possible imbalance between the number of agents and objects.

THEOREM 1. *Suppose any round of TTC begins with n agents and o objects remaining in the market. Then, the probability that there are $m \leq \min\{o, n\}$ agents assigned at the end of that*

[17] By reporting truthfully about one's preferences, each agent obtains a lottery that stochastically dominates another lottery he could obtain by misreporting his preferences.

round is

$$p_{n,o;m} = \left(\frac{m}{(on)^{m+1}}\right)\left(\frac{n!}{(n-m)!}\right)\left(\frac{o!}{(o-m)!}\right)(o+n-m).$$

Thus, denoting n_i and o_i the number of individuals and objects remaining in the market at any round i, the random sequence $\{(n_1,o_1),(n_2,o_2),...\}$ is a Markov chain.

Proof. See Appendix A. ∎

The Markov property stated in Theorem 1 is neither obvious nor anticipated, for the TTC algorithm in each round depends sensitively on the history of all preceding rounds. In particular, although the types (i.e., preferences and priorities) of agents are random in the first round, the types of the agents remaining in the subsequent rounds are no longer random, for they remain for "reasons." In particular, the "selection" of the remaining types makes them non-uniform,[18] and this makes the probabilistic method difficult to apply.[19]

The crucial analysis is to study the behavior of the so-called *random spanning forests* induced by TTC. The spanning forest at each round consists of vertices with outgoing edges of degree one and the vertices without any outgoing edges ("roots"). The former vertices are the agents and objects that pointed in the last round to those that did not form cycles and thus remain in the current round; clearly they must continue to point to the same parties in the current round. The roots are the agents and objects that pointed in the previous round to those that formed cycles and were thus removed at the end of that round; they must now randomly "repoint" to new parties. New cycles would emerge from this random repointing. It is easy to see that the resulting sequence of random spanning forests follows a Markov chain. The main analysis, and this comprises the bulk of the proof, is to prove that this Markov chain, when projected to the numbers n of agents and o of objects at each round, gives rise to a simpler Markov chain in (n,o). In addition, we compute transition probabilities, as stated in Theorem 1. Given the closed form solution for transition probabilities, we can further derive the expectation and the variance of the number of agents matched at each round (see Supplementary Appendix S.2).

[18] To see this more precisely, suppose there are three individuals, $1,2$ and 3, and three objects, o_1, o_2 and o_3. Per our assumption, the joint distribution of agents' preferences and priorities is uniform iid. Now consider a possible history say where agent 1 is assigned object o_1 in round 1 of TTC. At the beginning of round 2, the joint distribution of the preferences and priorities of the *remaining* agents 2 and 3 with regard to the *remaining* objects o_2 and o_3 is no longer uniform. Specifically, the probability that agent 2 prefers o_2 to o_3 is strictly less than $1/2$ *conditional* on her having a higher priority than agent 3 at o_2; the simple intuition is that the mismatched pointing between 2 and o_2 is weighted heavily in the conditioning as a possible reason for their remaining in the market. The precise calculation is available from the authors.

[19] In particular, the *principle of deferred decisions*—a standard method which views each agent as *drawing* preferences of the remaining objects at random in each round instead of having drawn preferences for objects in the beginning—does not apply in our context.

These characterizations prove useful for studying the *completion time* of TTC, i.e., the number of rounds it takes for the TTC algorithm to complete. Specifically, our characterization implies that if there remain n agents and n objects at the beginning of a round, an order $\sqrt{n}$ of agents and objects are matched on average in that round. Using this fact, we can show that the completion time for TTC is sublinear in n.[20,21]

PROPOSITION 1. *Let T denote the number of rounds required for TTC to conclude. Then,* $\frac{T}{n} \xrightarrow{p} 0$.

PROOF. See Supplementary Material S.5. ∎

The sublinear completion time is important for our purpose since we can later bound the expected number of short cycles per round. Combined with this bound, the sublinear completion time will eventually imply that short cycles become "rare" as the market grows large.

5 The Asymptotic Irrelevance of Priorities in TTC

We are now in a position to establish the asymptotic irrelevance of priorities in TTC. For this purpose, we consider a sequence of random markets with n agents and n objects. We then show that the **normalized priority ranks**—*the priority ranks divided by n*—enjoyed by all objects under TTC converge to uniform distribution over $[0,1]^n$, the outcome under RSD, independently of the preference ranks enjoyed by all agents.

To this end, we begin with the observation that if an object is assigned via a long cycle, it is not informative about the priority of the agent obtaining that object. To be more precise, fix an object o and let R_o and R_o^* denote respectively the priority rank at o of the agent *who obtains o* and the priority rank of the agent *whom o points to* at the round that it is assigned,

[20] The argument is roughly as follows. The fact that an order $\sqrt{n}$ of agents are cleared on average starting from n agents and objects means that with strictly positive probability (bounded away from zero), a very small number of objects (less than δn for an arbitrariliy small $\delta > 0$) remain after order $\sqrt{n}$ of rounds, as $n \to \infty$. This result further implies that with probability arbitrarily close to one the completion time is less than αn, for any $\alpha > 0$ arbitrarily small, as $n \to \infty$.

[21] It is known that the Shapley-Scarf TTC takes about $\sqrt{n}$ rounds (more precisely, $\sqrt{\frac{8}{\pi}n - \frac{3}{\pi}\log(n) + O(1)}$) to complete (see Theorem 1 of Frieze and Pittel (1995)). In general, one would expect our TTC to take longer to complete than Shapley-Scarf TTC. The reason is that, unlike our TTC, distinct objects always point to distinct agents in each round of Shapley-Scarf TTC (due to the bijective ownership structure), and this will lead to more cycles being formed in each round. Indeed, we show in the Supplementary Appendix S.3 that the number of rounds required for our TTC stochastically dominates the number of rounds required for Shapley-Scarf TTC. Nevertheless, Proposition 1 shows that, while slower than Shapley-Scarf TTC, the completion time of (prioritized) TTC is still sublinear.

respectively. Then, we can show that R_o is uniform random among $\{R_o^*+1,...,n\}$ conditional on $R_o > R_o^*$ (i.e., o being assigned via a long cycle). This follows from the fact that if an object o is assigned via a long cycle, then permuting the priority ranks at o of all agents remaining at that round (except for the agent o points to), including the agent receiving o, does not change the outcome of TTC. Since every such permutation is equally likely, the assignment of object o is uniform over remaining agents; and so his priority rank is uniform over $\{R_o^*+1,...,n\}$. The formal argument is provided in the Supplementary Appendix S.6.1.

Given this observation, our irrelevance result would follow if one shows that (i) TTC allocates virtually all objects via long cycles rather than short cycles and that (ii) for virtually all o, the $R_o^*/n \to 0$, so that the normalized priority rank R_o/n converges to $U[0,1]$, just like RSD, for almost all objects assigned via long cycles.

Formally, let $\hat{O} := \{o \in O | R_o > R_o^*\}$ denote the set of all objects assigned via long cycles and let

$$\tilde{O} := \{o \in O | R_o^* \leq \log^{1+\varepsilon}(n)\},$$

denote the set of objects that point to top $\log^{1+\varepsilon}(n)$-ranked agent at the time of assignment (where ε is an arbitrary strictly positive number). We first prove that their intersection $\bar{O} := \hat{O} \cap \tilde{O}$ eventually comprises the entire proportion of the objects as $n \to \infty$.

Lemma 1. *$\frac{|\bar{O}|}{n} \xrightarrow{p} 1$ as $n \to \infty$.*

Proof. It suffices to prove that $\frac{|\hat{O}|}{n} \xrightarrow{p} 1$ and $\frac{|\tilde{O}|}{n} \xrightarrow{p} 1$. We first prove the former. Corollary S1 of our Supplementary Appendix proves that at each round of TTC, irrespective of the history, the expected number of objects matched via short cycles is smaller than 2. Thus, denoting by $\hat{o}_t$ the number of objects involved in short cycles at Step t of TTC, we must have $\mathbb{E}[\hat{o}_t \,|T = t'] \leq 2$ for any $t \leq t'$. Hence,

$$\begin{aligned} \frac{1}{n}\mathbb{E}[|\hat{O}|] &= 1 - \frac{1}{n}\mathbb{E}\left[\sum_{t=1}^{T} \hat{o}_t\right] \\ &= 1 - \frac{1}{n}\mathbb{E}_T\left[\mathbb{E}\left[\sum_{t=1}^{T} \hat{o}_t \,|T\right]\right] \\ &\geq 1 - 2\mathbb{E}_T\left[\frac{T}{n}\right] \to 1, \end{aligned}$$

where the convergence result comes from Proposition 1. Since $\frac{1}{n}\mathbb{E}[|\hat{O}|] \to 1$ implies $\frac{|\hat{O}|}{n} \xrightarrow{p} 1$, we are done.

We next prove that $\frac{|\tilde{O}|}{n} \xrightarrow{p} 1$. To this end, we define a new mechanism TTC*, which operates exactly like TTC, except that, in each round, objects in each cycle are assigned to the agents that *the objects point to* (rather than the other way around). Clearly, in each round of TTC*, the same cycles as those in the corresponding round of TTC are formed, and the same associated set of agents and objects are assigned and removed. One crucial difference, though, is that the assignment is Pareto efficient *from the perspective of objects*. Proposition 1 of Che and Tercieux (2018), applied to the object side, then implies the result. ∎

We are almost ready for the main theorem. In the sequel, if W is a random variable defined on $\{1, ..., n\}$ then we let $\bar{W} := \frac{1}{n}W$. Denote by $O^n := \{o_1, ..., o_n\}$ and $I^n := \{i_1, ..., i_n\}$ respectively the set of objects and the set of individuals in an n-economy. Let $\{V_{o_1}, ..., V_{o_n}, V_{i_1}, ..., V_{i_n}\}$ be a sequence of collections of $2n$ random variables, each taking values in $[0, 1]$. We need to define a notion of convergence for a (random) vector whose length increases as $n \to \infty$:

Definition 1. *A random vector $\{\bar{W}_{o_1}, ..., \bar{W}_{o_n}, \bar{W}_{i_1}, ..., \bar{W}_{i_n}\}$ **converges in distribution to** $\{V_{o_1}, ..., V_{o_n}, V_{i_1}, ..., V_{i_n}\}$ as $n \to \infty$ if for any integer K, any $x \in [0, 1]^K$ and any sequence $\{y^n\}$ with values in $[0, 1]^n$, we have*

$$\lim_{n\to\infty} |F^n(x, y^n) - G^n(x, y^n)| = 0$$

where F^n is the cdf of $\{\bar{W}_{o_1}, ..., \bar{W}_{o_K}, \bar{W}_{i_1}, ..., \bar{W}_{i_n}\}$ while G^n is the cdf of $\{V_{o_1}, ..., V_{o_K}, V_{i_1}, ..., V_{i_n}\}$.

From now on, let $\{\bar{R}_{o_1}, ..., \bar{R}_{o_n}, \bar{R}_{i_1}, ..., \bar{R}_{i_n}\}$ denote the normalized priority and preference ranks under TTC. Since RSD ignores priorities, Pathak and Sethuraman (2011)'s equivalence result implies that the distribution of (normalized) ranks enjoyed by objects and agents under RSD is exactly $\{\bar{U}_{o_1}, ..., \bar{U}_{o_n}, \bar{R}_{i_1}, ..., \bar{R}_{i_n}\}$, where $\bar{U}_{o_1}, ..., \bar{U}_{o_n}$ denotes a collection of iid random variables each being $U\{1, ..., n\}/n$.

Recall our observation that any object assigned via long cycles is uniform-randomly assigned across individuals whom the object ranks below R^*_o (i.e., ranks larger than R^*_o). This in turn means that the rank enjoyed by each object $o \in \bar{O}$ is stochastically dominated by the uniform distribution across $\{\lceil \log^{1+\varepsilon}(n) \rceil + 1, ..., n\}$. In addition, the rank enjoyed by each object in $\bar{O}$ stochastically dominates the uniform distribution from $\{1, ..., n\}$. This is true independently of the distribution of the ranks enjoyed by the agents and the ranks enjoyed by the other objects in the set $\bar{O}$. This is proved in Proposition S2 of the supplementary material. Since $\log^{1+\varepsilon}(n)/n \to 0$ as $n \to \infty$, this proposition and Lemma 1 imply that the joint distribution of (normalized) ranks converges to the iid uniform distribution as $n \to \infty$.

The main theorem now follows: under TTC, the limit distribution of ranks enjoyed by the objects as $n \to \infty$ is uniform just like RSD.

Theorem 2 (Irrelevance). *(a) The profile $\{\bar{R}_{o_1}, ..., \bar{R}_{o_n}, \bar{R}_{i_1}, ..., \bar{R}_{i_n}\}$ of normalized priority and preference ranks under TTC converges in distribution to $\{\bar{U}_{o_1}, ..., \bar{U}_{o_n}, \bar{R}_{i_1}, ..., \bar{R}_{i_n}\}$ as $n \to \infty$;*

(b) For any $x \in [0,1]$, $\frac{1}{n}\sum_{o\in O}\mathbf{1}_{\{\bar{R}_o \leq x\}} \xrightarrow{p} x$.

Part (a) strengthens the well-known equivalence between TTC and RSD to the *joint* distribution of preference and priority: the joint distribution under TTC converges in distribution to the joint distribution of ranks of RSD. Part (b) establishes irrelevance of priorities in TTC, showing that the empirical distribution of the priority ranks under TTC converges in distribution to a collection of iid uniform random variables over $[0, 1]$.

Our result also has natural implications when comparing TTC and RSD in terms of justified envy. Indeed, the next result shows that relative to the size of the market, the incidences of justified envy under TTC and RSD become indistinguishable.

Corollary 1. *Fix a pair* (i, o). *The probability that* i *has a higher priority for* o *than its recipient under TTC converges to* $1/2$*—the same as that under RSD. Consequently, the difference between the probability that* (i, o) *blocks TTC and the probability that* (i, o) *blocks RSD goes to* 0*, and the difference in expected fraction of blocking pairs under TTC and RSD converges to* 0.[22]

Proof. See Appendix S.7. ∎

6 Discussion

The preceding irrelevance result suggests that the ability by TTC to use priorities to eliminate justified envy is increasingly limited as the market grows large. However, one may not interpret this as a design flaw of TTC. Recall from Abdulkadiroglu, Che, Pathak, Roth, and Tercieux (2017) that TTC is justified envy minimal. In light of this result, the vanishing role of priorities should instead be interpreted as a consequence of Pareto efficiency and strategyproofness; namely, these two requirements prove too strong to leave any scope for priorities to play a significant role in a large market.[23] Also, as we argue below, there are circumstances in which the role of priorities may not vanish even in a large market, although our asymptotic irrelevance result is robust in several ways.

6.1 Correlated preferences:

If agents' preferences are perfectly correlated, then TTC completely eliminates justified envy. Plainly, all objects would be assigned via short cycles in this case, so the resulting assignment

[22]The fraction of blocking pair corresponds to the total number of blocking pairs divided by the total number of possible pairs $|I| \times |O|$.

[23]Indeed, in our environment, one can find a non-strategyproof Pareto-efficient mechanism, e.g., Boston mechanism, that admits significantly less justified envy than RSD even in a large market.

would be stable (as well as efficient). By contrast, RSD would entail a significant amount of justified envy even in the limit. But this is an extreme case that is unlikely to hold in any real setting. A more interesting case is when the preferences are not perfectly correlated. Suppose for instance that agents' preferences are represented by a cardinal utility function, $u_i(o) = u_o + \xi_{io}$, where u_o is a common utility from object o common for all agents and ξ_{io} is an idiosyncratic utility from o drawn iid for each agent i. Suppose further the support of both shocks are bounded. Our asymptotic irrelevance result appears to hold in this environment.[24] For instance, Figure 2 shows that asymptotic irrelevance extends to the case of correlated preferences where the u_o and ξ_{io} are each distributed from $U[0, 1]$.

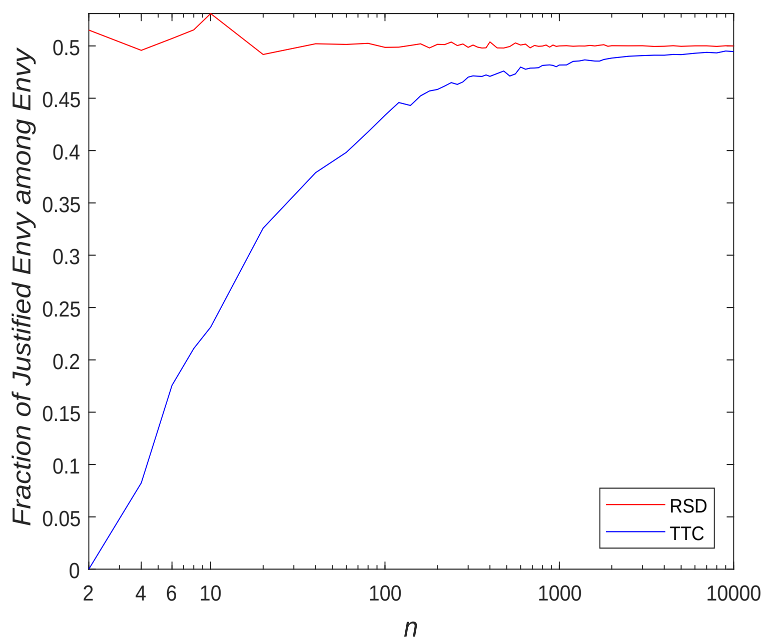

Figure 2: Incidences of justified envy among incidences of envy.

Note: The horizontal axis indicates the number of agents/objects. The figure plots the average proportions across 100 draws of $(u_o, \xi_{io}, \eta_{io})$ from $U[0, 1]^3$, where η_{io} is priority score of agent i at o.

[24] To see this analytically, consider a simpler case in which the support of common utility is finite and sufficiently far apart from each other that the objects are effectively "tiered": all agents prefer top tier objects (with the highest value of u_o), and they all prefer the second tier objects next, and so on. In this case, the TTC are effectively partitioned into multiple stages: in stage 1, all agents point to objects in tier 1, and once all tier 1 objects are assigned, stage 2 begins in which the remaining agents point to tier 2 objects, and they are assigned, etc. Since agents' priorities at each object is iid, each stage can be separated as a distinct TTC market, for which our asymptotic irrelevance result would apply.

6.2 Many-to-one matching:

Our model has considered an one-to-one matching environment. While one-to-one matching serves as a good baseline model, many real-world situations involve many-to-one matchings. School choice or housing allocation typically involves multiple seats or multiple identical units available for assignment. Our result of asymptotic irrelevance appears to hold for many-to-one matching setting as long as the number of copies per object type grows sufficiently slowly compared with the number of object types.[25] This can be seen in Figure 3 below, where the number of object types increases linearly with n, the number of agents, while the number of copies per object type is fixed at 20. (Preferences and priorities are still drawn iid uniformly.)

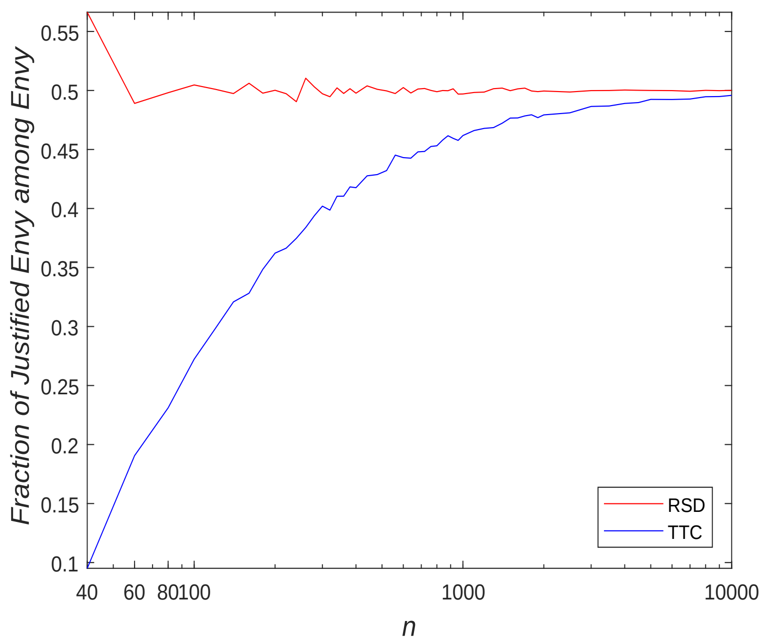

Figure 3: Incidences of justified envy among incidences of envy.

Note: The horizontal axis indicates the number n of agents. The number of objects is equal to $n/20$. The figures plot averages over 100 (uniform iid) draws of agents' preferences and objects' priorities.

At the same time, there is a sense in which the irrelevance result does not extend to the other commonly used asymptotics, in which the number of copies per object type grows

[25]Such asymptotics, which one may call a "small school" model, has been adopted by a number of authors, such as Kojima and Pathak (2009), Ashlagi, Kanoria, and Leshno (2017), Che and Tercieux (2018) and Che and Tercieux (2015), fits well settings such as medical matching (where about 20,000 doctors apply to about 3,000-4,000 hospitals), and NYC public high school matching in which the number of programs (about 800) exceeds the number of students admitted by each program (about 100).

sufficiently fast compared with the number of object types.[26] To see this, suppose the number of object types is fixed at some finite number but the number of agents and the number of copies for each object type grows large. In that case, the proportion of agents that are assigned via short cycles under TTC does not vanish even in probability. Indeed, Panel (a) of Figure 4 shows that the differences between the two mechanisms do not shrink as the market grows large, in terms of the total incidences of justified envy. Our asymptotic result *does* extend, however, if one considers alternative measures, such as the number of agents with justified envy among those with envy.[27] Panel (b) of Figure 4 illustrates that the convergence holds.[28]

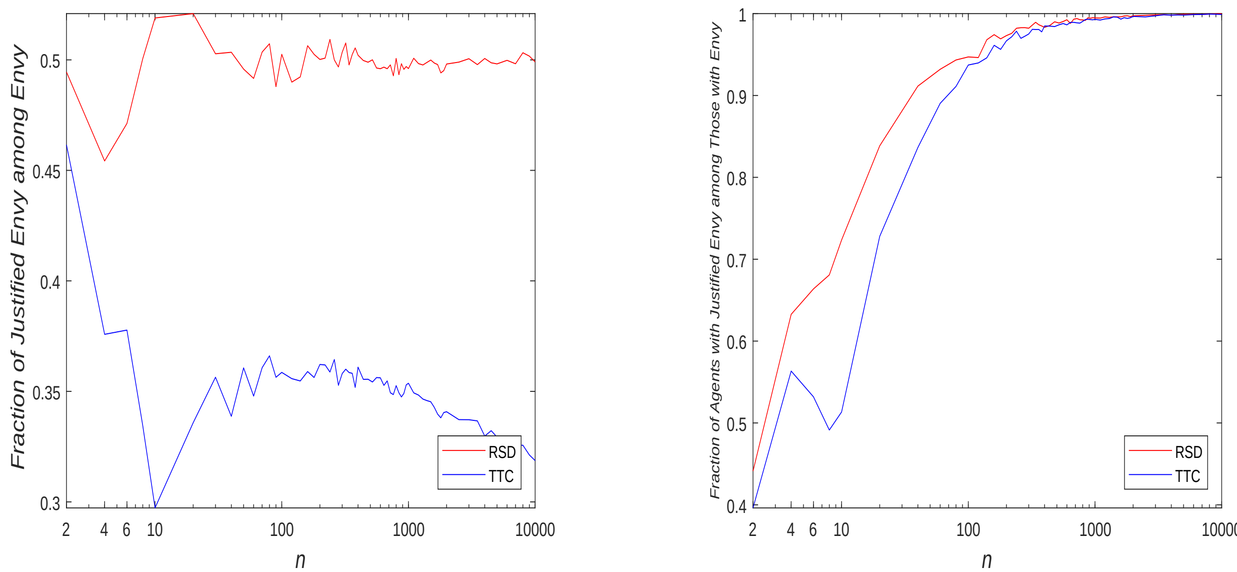

(a) Incidences of justified envy among incidences of envy. (b) Agents with justified envy among those with envy.

Figure 4: Difference between TTC and RSD

Note: The horizontal axis indicates the size of the market, i.e., the number of agents. The number of objects is fixed to 10. The figures plot averages over 100 (uniform iid) draws of agents' preferences and objects' priorities.

Ultimately, our contribution clarifies the circumstances in which the effect of priorities in TTC disappears and those in which TTC can be expected to perform better than RSD.

[26] Such asymptotics, which one may call a "large school" model, has been adopted by many authors such as Abdulkadiroglu, Che, and Yasuda (2015), Azevedo and Leshno (2016), Che, Kim, and Kojima (2013) and Leshno and Lo (2017), and fits well with the school choice in many US cities in which a handful of schools admit each hundreds of students.

[27] More precisely, we count the number of agents with justified envy and divide this by the number of agents with envy (not necessarily justifiable). In short, the measure (a) reflects the intensity of justified envy by counting all incidences of justified envy, whereas (b) only counts the number of agents who have justified envy, irrespective of how many agents they justifiably envy.

[28] Our convergence also holds for the fraction of blocking pairs.

References

Abdulkadiroglu, A., Y.-K. Che, P. Pathak, A. Roth, and O. Tercieux (2017): "Minimizing Justified Envy in School Choice: The Design of New Orleans' OneApp," Columbia, Duke, MIT, PSE, Stanford; Unpublished mimeo.

Abdulkadiroglu, A., Y.-K. Che, and Y. Yasuda (2015): "Expanding 'Choice' in School Choice," *American Economic Journal: Microeconomics*, 7, 1–42.

Abdulkadiroglu, A., P. A. Pathak, and A. E. Roth (2009): "Strategy-proofness versus Efficiency in Matching with Indifferences: Redesigning the NYC High School Match," *American Economic Review*, 99(5), 1954–1978.

Abdulkadiroglu, A., and T. Sönmez (1998): "Random Serial Dictatorship and the Core from Random Endowments in House Allocation Problems," *Econometrica*, 66, 689–701.

Abdulkadiroglu, A., and T. Sonmez (2003): "School Choice: A Mechanism Design Approach," *American Economic Review*, 93, 729–747.

Ashlagi, I., Y. Kanoria, and J. D. Leshno (2017): "Unbalanced Random Matching Markets," *Journal of Political Economy*, 125, 69–98.

Azevedo, E. M., and J. D. Leshno (2016): "A supply and demand framework for two-sided matching markets," *Journal of Political Economy*, 124, 1235–1268.

Bade, S. (2016): "Random Serial Dictatorship: the One and Only," Royal Holloway College, Unpublished mimeo.

Balinski, M., and T. Sönmez (1999): "A tale of two mechanisms: student placement," *Journal of Economic Theory*, 84, 73–94.

Carroll, G. (2014): "A General Equivalence Theorem for Allocation of Indivisible Objects," *Journal of Mathematical Economics*, 51, 163–177.

Che, Y.-K., J. Kim, and F. Kojima (2013): "Stable Matching in Large Economies," mimeo.

Che, Y.-K., and O. Tercieux (2015): "Efficiency and Stability in Large Matching Markets," forthcoming, *Journal of Political Economy*.

——— (2018): "Payoff Equivalence of Efficient Mechanisms in Large Markets," .

Frieze, A., and B. Pittel (1995): "Probabilistic Analysis of an Algorithm in the Theory of Markets in Indivisible Goods," *The Annals of Applied Probability*, 5, 768–808.

Gale, D., and L. S. Shapley (1962): "College Admissions and the Stability of Marriage," *American Mathematical Monthly*, 69, 9–15.

Knuth, D. E. (1996): "An Exact Analysis of Stable Allocation," *Journal of Algorithms*, 2, 431–444.

Kojima, F., and P. A. Pathak (2009): "Incentives and Stability in Large Two-Sided Matching Markets," *American Economic Review*, 99, 608–627.

Leshno, J. D., and I. Lo (2017): "The Cutoff Structure of Top Trading Cycles in School Choice," Columbia University, Unpublished mimeo.

Pathak, P., and J. Sethuraman (2011): "Lotteries in Student Assignment: An Equivalence Result," *Theortical Economics*, 6, 1–17.

Pathak, P. A., and J. Sethuraman (2010): "Lotteries in Student Assignment: An Equivalence Result," *Theoretical Economics*.

Roth, A. E. (1982): "The Economics of Matching: Stability and Incentives," *Mathematics of Operations Research*, 7, 617–628.

Shapley, L., and H. Scarf (1974): "On Cores and Indivisibility," *Journal of Mathematical Economics*, 1, 22–37.

Sonmez, T., and U. Unver (2011): "Market Design for Kidney Exchange," in *Oxford Handbook of Market Design*, ed. by M. N. Z. Neeman, and N. Vulkan. forthcoming.

A Proof of Theorem 1

Given the nature of conditioning mentioned earlier, it is crucial for our purpose to keep track of the agents and objects that *can* draw their partners at random and those who *cannot* in each round of TTC. This requires us to investigate the probabilistic structure known as random rooted forests.

To begin, consider any two finite sets I and O, with cardinalities $|I| = n, |O| = o$. A **bipartite digragh** $G = (I \cup O, E)$ consists of vertices I and O on two separate sides and directed edges $E \subset (I \times O) \cup (O \times I)$, comprising ordered pairs of the form (i, o) or (o, i) (corresponding to edge originating from i and pointing to o and an edge from o to i, respectively). A **rooted tree** is a bipartite digraph where all vertices have out-degree 1 except the root which has out-degree 0.[29] A **rooted forest** is a bipartite graph which consists of a collection of disjoint rooted trees. A **spanning rooted forest over** $I \cup O$ is a forest comprising vertices $I \cup O$. From now on, a spanning forest will be understood as being over $I \cup O$.

A.1 Markov Properties of Spanning Rooted Forests Induced by TTC

We begin by noting that TTC induces a random sequence of spanning rooted forests. Indeed, one could see the beginning of the first round of TTC as a situation where we have the trivial forest consisting of $|I|+|O|$ roots, or isolated vertices without edges. In the example of Section 4, there are 6 separate trees: $\{1\}, \{2\}, \{3\}, \{a\}, \{b\}, \{c\}$. Within this step, each vertex in I randomly points to a vertex in O and each vertex in O randomly points to a vertex in I. Assume that agents 1 and 3 point to c, and agent 2 points to a, and all objects point to 3. Note that once we delete the realized cycles ($3 - c$ in the example), we again get a spanning rooted forest. So we can again view the graph at the beginning of the second round of TTC as a spanning rooted forest, where the roots consist of those agents and objects that had pointed to the entities that were cleared via cycles. In the above example, the spanning rooted forest in the beginning of Round 2 has three rooted trees: $\{1\}, \{2 \to b\}, \{a\}$. Here again objects that are roots randomly point to a remaining individual and individuals that are roots randomly point to a remaining object. Once cycles are cleared we again obtain a forest and the process goes on like this.

Formally, the random sequence of forests, $F_1, F_2, ...$, is defined as follows. First, we let F_1 be a trivial unique forest consisting of $|I| + |O|$ trees with isolated vertices, forming their own roots. For any $i = 2, ...$, we first create a random directed edge from each root of F_{i-1} to a vertex on the other side, and then delete the resulting cycles (these are the agents and objects assigned in round $i - 1$) and F_i is defined to be the resulting rooted forest. Note that this random sequence of spanning rooted forests is a Markov chain.

For any rooted forest F_i, let $N_i = I_i \cup O_i$ be its vertex set and $k_i = (k_i^I, k_i^O)$ be the vector denoting the numbers of roots on both sides, and use (N_i, k_i) to summarize this information. And let $\mathcal{F}_{N_i,k_i}$ denote the set of all rooted forests having N_i as the vertex set and k_i as the vector of its root numbers.

Given N_i, k_i and N_{i+1}, k_{i+1}, for each forest $F \in \mathcal{F}_{N_{i+1},k_{i+1}}$, one can compute the number of possible pairs (F', ϕ) that could have given rise to F, where $F' \in \mathcal{F}_{N_i,k_i}$ and ϕ maps the roots

[29]Sometimes, a tree is defined as an acyclic undirected connected graph. In such a case, a tree is rooted when we name one of its vertex a "root." Starting from such a rooted tree, if all edges now have a direction leading toward the root, then the out-degree of any vertex (except the root) is 1. So the two definitions are actually equivalent.

of F' in I_i to its vertices in O_i as well as the roots of F' in O_i to its vertices in I_i. In words, such a pair (F', ϕ) corresponds to a set N_i of agents and objects remaining at the beginning of round i of TTC, of which k_i^I agents of I_i and k_i^O objects have lost their favorite parties (and thus they must repoint to new partners in N_i under TTC in round i), and the way in which they repoint to the new partners under TTC in round i causes the new forest F to emerge at the beginning of round $i+1$ of TTC (after cycles are cleared). There are typically multiple such pairs (F', ϕ) that could give rise to F. Let $\beta(I_i, O_i, k_i^I, k_i^O; I_{i+1}, O_{i+1}, k_{i+1}^I, k_{i+1}^O)$ denote the number of pairs (F', ϕ), $F' \in \mathcal{F}_{N_i,k_i}$, causing F to arise.

The first important observation we make is that β does not depend on the particular $F \in \mathcal{F}_{N_{i+1},k_{i+1}}$. That is, its dependence on sets $(I_i, O_i, ; I_{i+1}, O_{i+1})$ is only through their cardinalities—specifically, $|I_i|$, $|O_i|$, k_i, k_{i+1}, $|I_{i+1}|$, and $|O_{i+1}|$, or equivalently $|I_i|$, $|O_i|$, k_i, k_{i+1}, and $|N_i| - |N_{i+1}|$.

In order to show that this number does not depend on the particular $F \in \mathcal{F}_{N_{i+1},k_{i+1}}$, let us, for any given $F \in \mathcal{F}_{N_{i+1},k_{i+1}}$, construct all such pairs. To build any such pair (F', ϕ), one can simply choose a quadruplet (a, b, c, d) of four non-negative integers with $a + c = k_i^I$ and $b + d = k_i^O$ and proceed by,

(i) choosing c old roots from I_{i+1}, and similarly, d old roots from O_{i+1},

(ii) choosing a old roots from $I_i \backslash I_{i+1}$ and similarly, b old roots from $O_i \backslash O_{i+1}$,

(iii) choosing a partition into cycles of $N_i \backslash N_{i+1}$, each cycle of which contains at least one old root from (ii),[30]

(iv) choosing a mapping of the $k_{i+1}^I + k_{i+1}^O$ new roots to $N_i \backslash N_{i+1}$.[31]

Clearly, the number of pairs (F', ϕ), $F' \in \mathcal{F}_{N_i,k_i}$, satisfying the above restrictions depends only on $|I_i|$, $|O_i|$, k_i, k_{i+1}, and $|N_i| - |N_{i+1}|$.[32] Accordingly, we now write this number as $\beta(|I_i|, |O_i|, k_i; |N_i| - |N_{i+1}|, k_{i+1})$. (We derive β explicitly in Section S.1.1 of the supplementary material.)

With this observation in hand, we can prove the following result.

LEMMA 2. *Given $(N_j, k_j), j = 1, ..., i$, every (rooted) forest of $\mathcal{F}_{N_i,k_i}$ is equally likely.*

[30]Within round i of TTC, one cannot have a cycle creating only with vertices that are not roots in the forest obtained at the beginning of round i. This is due to the simple fact that a forest is an acyclic graph. Thus, each cycle creating must contain at least one old root. Given that, by definition, these roots are eliminated from the set of available nodes in round $i + 1$, these old roots that each cycle must contain must be from (ii).

[31]Since, by definition, any root in $F \in \mathcal{F}_{N_{i+1},k_{i+1}}$ does not point, this means that, in the previous round, this node was pointing to another node which was eliminated at the end of that round.

[32]Recall that by definition of TTC, whenever a cycle creates, the same number of individuals and objects must be eliminated in this cycle. Hence, $|O_i| - |O_{i+1}| = |I_i| - |I_{i+1}|$ and $|N_i| - |N_{i+1}| = 2(|I_i| - |I_{i+1}|)$.

PROOF. We prove this result by induction on i. Since for $i = 1$, by construction, the trivial forest is the unique forest which can occur, this is trivially true for $i = 1$. Fix $i \geq 1$, and assume our statement is true for i. Let us show that it holds for $i + 1$.

Fix $N_i = I_i \cup O_i \supset N_{i+1} = I_{i+1} \cup O_{i+1}$, and k_i and k_{i+1}. For each forest $F \in \mathcal{F}_{N_{i+1},k_{i+1}}$, we consider a possible pair (F', ϕ) that could have given rise to F, where $F' \in \mathcal{F}_{N_i,k_i}$ and ϕ maps the roots of F' in I_i to its vertices in O_i as well as the roots of F' in O_i to its vertices in I_i. As we already argued, each forest $F \in \mathcal{F}_{N_{i+1},k_{i+1}}$ arises from the same number of such pairs—i.e., that the number of pairs (F', ϕ), $F' \in \mathcal{F}_{N_i,k_i}$, causing F to arise does not depend on the particular $F \in \mathcal{F}_{N_{i+1},k_{i+1}}$.

The number of such pairs is given by $\beta(|I_i|, |O_i|, k_i; |N_i| - |N_{i+1}|, k_{i+1})$. Let $\phi_i = (\phi_i^I, \phi_i^O)$ where ϕ_i^I is the random mapping from the roots of F_i in I_i to O_i and ϕ_i^O is the random mapping from the roots of F_i in O_i to I_i. Let $\phi = (\phi^I, \phi^O)$ be a generic mapping of that sort. Since, conditional on $F_i = F'$, the mappings ϕ_i^I and ϕ_i^O are uniform, we get

$$\Pr(F_{i+1} = F | F_i = F') = \frac{1}{|O_i|^{k_i^I}} \frac{1}{|I_i|^{k_i^O}} \sum_{\phi} \Pr(F_{i+1} = F | F_i = F', \phi_i = \phi). \tag{1}$$

Therefore, we obtain

$$\begin{aligned}
&\Pr(F_{i+1} = F | (N_1, k_1), ..., (N_i, k_i)) \\
=& \sum_{F' \in \mathcal{F}_{N_i,k_i}} \Pr(F_{i+1} = F, F_i = F' | (N_1, k_1), ..., (N_i, k_i)) \\
=& \sum_{F' \in \mathcal{F}_{N_i,k_i}} \Pr(F_{i+1} = F | (N_1, k_1), ..., (N_i, k_i), F_i = F') \Pr(F_i = F' | (N_1, k_1), ..., (N_i, k_i)) \\
=& \frac{1}{|\mathcal{F}_{N_i,k_i}|} \sum_{F' \in \mathcal{F}_{N_i,k_i}} \Pr(F_{i+1} = F | F_i = F') \\
=& \frac{1}{|\mathcal{F}_{N_i,k_i}|} \sum_{F' \in \mathcal{F}_{N_i,k_i}} \frac{1}{|O_i|^{k_i^I}} \frac{1}{|I_i|^{k_i^O}} \sum_{\phi} \Pr(F_{i+1} = F | F_i = F', \phi_i = \phi) \\
=& \frac{1}{|\mathcal{F}_{N_i,k_i}|} \frac{1}{|O_i|^{k_i^I}} \frac{1}{|I_i|^{k_i^O}} \sum_{F' \in \mathcal{F}_{N_i,k_i}} \sum_{\phi} \Pr(F_{i+1} = F | F_i = F', \phi_i = \phi) \\
=& \frac{1}{|\mathcal{F}_{N_i,k_i}|} \frac{1}{|O_i|^{k_i^I}} \frac{1}{|I_i|^{k_i^O}} \beta(|I_i|, |O_i|, k_i; |N_i| - |N_{i+1}|, k_{i+1}),
\end{aligned} \tag{2}$$

where the third equality follows from the Markov property of $\{F_j\}$ and the induction hypothesis, the fourth follows from (1), and the last follows from the definition of β and from the fact that the conditional probability in the sum of the penultimate line is 1 or 0, depending upon whether the forest F arises from the pair $(F^{'}, \phi)$ or not. Note that this probability is independent of $F \in \mathcal{F}_{N_{i+1},k_{i+1}}$. Hence,

$$\begin{aligned}
&\Pr(F_{i+1} = F|(N_1, k_1), ..., (N_i, k_i), (N_{i+1}, k_{i+1})) \\
=&\frac{\Pr(F_{i+1} = F|(N_1, k_1), ..., (N_i, k_i))}{\Pr(F_{i+1} \in \mathcal{F}_{N_{i+1},k_{i+1}}|(N_1, k_1), ..., (N_i, k_i))} \\
=&\frac{\Pr(F_{i+1} = F|(N_1, k_1), ..., (N_i, k_i))}{\sum_{\tilde{F}\in\mathcal{F}_{N_{i+1},k_{i+1}}} \Pr(F_{i+1} = \tilde{F}|(N_1, k_1), ..., (N_i, k_i))} = \frac{1}{|\mathcal{F}_{N_{i+1},k_{i+1}}|},
\end{aligned}$$

which proves that, given $(N_j, k_j), j = 1, ..., i + 1$, every rooted forest of $\mathcal{F}_{N_{i+1},k_{i+1}}$ is equally likely. ∎

The next lemma then follows easily.

LEMMA 3. *Random sequence (N_i, k_i) forms a Markov chain.*

PROOF. By (2) we must have

$$\begin{aligned}
\Pr((N_{i+1}, k_{i+1})|(N_1, k_1), ..., (N_i, k_i)) &= \sum_{F\in\mathcal{F}_{N_{i+1},k_{i+1}}} \Pr(F_{i+1} = F|(N_1, k_1), ..., (N_i, k_i)) \\
&= \sum_{F\in\mathcal{F}_{N_{i+1},k_{i+1}}} \frac{1}{|\mathcal{F}_{N_i,k_i}|} \frac{1}{|O_i|^{k_i^I}} \frac{1}{|I_i|^{k_i^O}} \beta(|I_i|, |O_i|, k_i; |N_{i+1}| - |N_i|, k_{i+1}).
\end{aligned}$$

Observing that the conditional probability depends only on (N_{i+1}, k_{i+1}) and (N_i, k_i), the Markov chain property is established. ∎

The proof of Lemma 3 reveals in fact that the conditional probability of (N_{i+1}, k_{i+1}) depends on N_i only through its cardinalities $(|I_i|, |O_i|)$, leading to the following conclusion. Let $n_i := |I_i|$ and $o_i := |O_i|$.

COROLLARY 2. *Random sequence $\{(n_i, o_i, k_i^I, k_i^O)\}$ forms a Markov chain.*

PROOF. See Supplementary Appendix S.1.3. ∎

A.2 Proof of Theorem 1

We first compute the probability of transition from (n_i, o_i, k_i^I, k_i^O) to $(n_{i+1}, o_{i+1}, k_{i+1}^I, k_{i+1}^O)$ such that $k_{i+1}^I = \lambda^I$ and $k_{i+1}^O = \lambda^O$:

$$\begin{aligned}
&\mathbf{P}(n, o, k_i^I, k_i^O; m, \lambda^I, \lambda^O) \\
:=&\Pr\left\{n_i - n_{i+1} = o_i - o_{i+1} = m, k_{i+1}^I = \lambda^I, k_{i+1}^O = \lambda^O \mid n_i = n, o_i = o, k_i^I, k_i^O\right\}.
\end{aligned}$$

This will be computed as a fraction $\frac{\Theta}{\Upsilon}$. The denominator Υ counts the number of rooted forests in the bipartite digraph with k_i^I roots in I_i and k_i^O roots in O_i, multiplied by the ways in which k_i^I roots of I_i could point to O_i and k_i^O roots of O_i could point to I_i.[33] We let $f(n, o, k_i^I, k_i^O)$ denote the number of rooted forests in a bipartite digraph (with n and o vertices on both sides) containing k_i^I and k_i^O roots on both sides. Thus, since there are $o^{k_i^I} n^{k_i^O}$ ways in which k_i^I roots in I_i point to O_i and k_i^O roots in O_i could point to I_i, $\Upsilon = o^{k_i^I} n^{k_i^O} f(n, o, k_i^I, k_i^O)$.

The numerator Θ counts the number of ways in which m agents are chosen from I_i and m objects are chosen from O_i to form a bipartite bijection each cycle of which contains at least one of $k_i^I + k_i^O$ old roots, and for each such choice, the number of ways in which the remaining vertices form a spanning rooted forest and the λ^I roots in I_{i+1} point to objects in $O_i \setminus O_{i+1}$ and λ^O roots in O_{i+1} point to agents in $I_i \setminus I_{i+1}$. To compute this, recall $\beta(n, o, k_i^I, k_i^O; m, \lambda^I, \lambda^O)$, denotes, for any F with $n-m$ agents and $o-m$ objects and roots λ^I and λ^O on both sides, the total number of pairs (F', ϕ) that could have given rise to F, where F' has n agents and o objects with (k_i^I, k_i^O) roots and ϕ maps the roots to the remaining vertices.

The numerator Θ is now computed as $\Theta = \binom{n}{m}\binom{o}{m} f(n-m, o-m, \lambda^I, \lambda^O)\beta(n, o, k_i^I, k_i^O; m, \lambda^I, \lambda^O)$.

In Section S.1.1 of the supplementary material, we explicitly compute $\beta(n, o, k_i^I, k_i^O; m, \lambda^I, \lambda^O)$, then, collecting terms, we show in the Supplementary Appendix S.1.4 that

$$\mathbf{P}(n, o, k_i^I, k_i^O; m, \lambda^I, \lambda^O) = \frac{\Theta}{\Upsilon} = \frac{1}{o^n n^o}\left(\frac{n!}{(n-m)!}\right)\left(\frac{o!}{(o-m)!}\right) m^{\lambda^I + \lambda^O} f(n-m, o-m, \lambda^I, \lambda^O).$$

A key observation is that this expression does not depend on (k_i^I, k_i^O), which, together with Corollary 2, implies that (n_i, o_i) forms a Markov chain.

Its transition probability can be derived by summing the expression over all possible (λ^I, λ^O)'s:

$$p_{n,o;m} := \sum_{0 \le \lambda^I \le n-m, 0 \le \lambda^O \le o-m} \mathbf{P}(n, o, k_i^I, k_i^O; m, \lambda^I, \lambda^O).$$

which is equal to the formula stated in the lemma $\left(\frac{m}{(on)^{m+1}}\right)\left(\frac{n!}{(n-m)!}\right)\left(\frac{o!}{(o-m)!}\right)(o+n-m)$ as shown in Supplementary Appendix S.1.4.

[33] Given that we have $n_i = n$ individuals, $o_i = o$ objects and k_i^I roots in I_i and k_i^O roots in O_i at the beginning of step i under TTC, one may think of this as the total number of possible bipartite digraph one may obtain via TTC at the end of step i when we let k_i^I roots in I_i point to their remaining most favorite object and k_i^O roots in O_i point to their remaining most favorite individual.