Efficiency, Justified Envy, and Incentives in Priority-Based Matching

AER: Insights 2020, 2(4): 425–442
https://doi.org/10.1257/aeri.20190307

Efficiency, Justified Envy, and Incentives in Priority-Based Matching†

By Atila Abdulkadiroğlu, Yeon-Koo Che, Parag A. Pathak, Alvin E. Roth, and Olivier Tercieux*

Top trading cycles (TTC) is Pareto efficient and strategy-proof in priority-based matching, but so are other mechanisms including serial dictatorship. We show that TTC minimizes justified envy among all Pareto-efficient and strategy-proof mechanisms in one-to-one matching. In many-to-one matching, TTC admits less justified envy than serial dictatorship in an average sense. Empirical evidence from New Orleans OneApp and Boston Public Schools shows that TTC has significantly less justified envy than serial dictatorship. (*JEL* C78, D61)

Important resources such as housing, organs, and school seats are allocated based on the participants' preferences and their priorities. Priorities reflect public policy or fairness considerations, such as seniority, the severity of needs or waiting list times, and geographic location or test scores. Two prominent mechanisms used for priority-based matching are Gale and Shapley's (1962) deferred acceptance (DA) and Gale's top trading cycles (TTC) (Shapley and Scarf 1974). Both mechanisms are strategy-proof: truthful reporting of preferences is a weakly dominant strategy for individuals.[1] DA eliminates justified envy; that is, no individual prefers another assignment over her assignment and has a higher priority than someone else assigned to the preferred assignment. TTC is Pareto efficient.

*Abdulkadiroğlu: Department of Economics, Duke University, and NBER (email: aa88@duke.edu); Che: Department of Economics, Columbia University (email: yc2271@columbia.edu); Pathak: Department of Economics, MIT, and NBER (email: ppathak@mit.edu); Roth: Department of Economics, Stanford University (email: alroth@stanford.edu); Tercieux: Paris School of Economics (email: tercieux@pse.ens.fr). Larry Samuelson was the coeditor for this article. We are grateful to Neil Dorosin, Gabriela Fighetti, and John White for assistance with this study. Thanks also to seminar participants at the 2016 ASSA meetings for helpful comments. Vira Semenova provided excellent research assistance. Pathak is grateful to the W. T. Grant Foundation and the National Science Foundation for support. Che is grateful to the National Science Foundation (SES-1851821) for support. This paper supersedes "The Role of Priorities in Assigning Indivisible Objects: A Characterization of Top Trading Cycles," cited by others as Abdulkadiroğlu, Atila, and Yeon-Koo Che (2010) or Abdulkadiroğlu, Atila, Yeon-Koo Che, and Olivier Tercieux (2010), and "Minimizing Justified Envy in School Choice: The Design of New Orleans' OneApp" (2017) by Abdulkadiroğlu, Atila, Yeon-Koo Che, Parag A. Pathak, Alvin E. Roth, and Olivier Tercieux. Roth is a member of the scientific advisory board of the Institute for Innovation in Public School Choice (IIPSC). IIPSC was involved in designing OneApp in New Orleans. Abdulkadiroğlu, Pathak, and Roth also advised Boston Public Schools and New York City's Department of Education on designing their student assignment systems, discussed herein. This article does not represent the views of the New Orleans Recovery School District or any other school district.
†Go to https://doi.org/10.1257/aeri.20190307 to visit the article page for additional materials and author disclosure statement(s).

[1] For original proofs for the one-to-one matching models, see Dubins and Freedman (1981), Roth (1982a), and Roth (1982b).

There is no mechanism that is both Pareto efficient and without justified envy.[2] Gale and Shapley (1962) showed that DA is constrained optimal since the DA matching weakly Pareto dominates any other matching without justified envy. In contrast, there are many Pareto-efficient and strategy-proof mechanisms. Does TTC provide a comparable constrained-optimal solution but with regard to elimination of justified envy?

This paper provides such a result for TTC in one-to-one matching. We define a partial order over efficient mechanisms by comparing instances of justified envy problem by problem. Specifically, mechanism φ_1 has less justified envy than mechanism φ_2 if, in each problem, every justified envy instance of φ_1 is also a justified envy instance of φ_2. That is, if an individual prefers another assignment over her φ_1 assignment and has a higher priority than someone else assigned to the preferred assignment, the individual also prefers another assignment over her φ_2 assignment and has a higher priority than someone else assigned to the preferred assignment. Our main result states that there is no Pareto-efficient and strategy-proof mechanism that has less justified envy than TTC in one-to-one matching. In other words, TTC is "justified envy minimal" according to set inclusion of justified envy instances within that class. However, our result does not mean that TTC is the unique such mechanism. Any two justified envy–minimal mechanisms are not comparable according to our partial order. This is in a similar vein with two Pareto-efficient mechanisms not being comparable with respect to the Pareto-domination relationship.

While TTC need not be the only justified envy–minimal mechanism, we formalize a sense in which TTC uses priorities correctly. Consider a class of mechanisms that are implemented by running TTC with some "artificial" priorities. If artificial priorities of such a mechanism differ from true priorities *even* for one individual at one object, the mechanism is no longer justified envy minimal. This class includes serial dictatorship (SD), a popular Pareto-efficient and strategy-proof mechanism.

Justified envy minimality of TTC does not extend to the many-to-one matching environment. We show, however, TTC outperforms SD—an obvious efficient alternative—by admitting less justified envy *in an average sense* when every possible priority profile is considered or when participants' priorities are drawn uniform randomly. Real-world data for priorities and preferences from New Orleans and Boston school assignment show that TTC has significantly fewer instances of justified envy compared to SD with uniformly drawn random serial orders.

The absence of a rigorous but simple description of the role of priorities in TTC may explain its limited use in practice. A 2005 Boston task force initially recommended TTC over DA for school assignment, stating the following (Landsmark, Dajer, and Gonsalves 2004):

> The Gale-Shapley algorithm…cuts down on the amount of choice afforded to families. The Top Trading Cycles algorithm also takes into account priorities while leaving some room for choice.…Choice was very important to many families who attended community forums.

[2] This trade-off is highlighted by Abdulkadiroğlu and Sönmez (2003), and it is a consequence of an example due to Roth (1982b).

This recommendation was eventually overturned for DA (Abdulkadiroğlu et al. 2005). The final report criticized TTC as follows (BPS Strategic Planning Team 2005):

> [TTC's] trading shifts the emphasis onto the priority and away from the goals BPS is trying to achieve by granting these priorities in the first place.

TTC was adopted in New Orleans as part of its OneApp assignment system in 2012 (Vanacore 2012). As far as we know, New Orleans is the only place that TTC has ever been used in practice. However, after one year, officials abandoned TTC for DA, in part due to its treatment of priorities. Our clarification of the role of priorities in TTC provides new ammunition for considering TTC in priority-based matching, particularly in the one-to-one case.

This paper is related to several on axiomatic mechanism design. Ma's (1994) TTC characterization for the housing market model is closest. The characterizations in Pápai (2000) and Pycia and Ünver (2017) are also related, but neither concerns justified envy. Dur (2013) and Morrill (2015) each characterize TTC by relying on additional axioms, which are distinct from our result.

I. The Model

A priority-based matching problem has the following ingredients:

(i) agents $I = \{i_1, \ldots, i_n\}$,

(ii) objects $S = \{s_1, \ldots, s_m\}$,

(iii) strict agent preferences $P = (P_{i_1}, \ldots, P_{i_n})$, and

(iv) strict object priorities $\succ = (\succ_{s_1}, \ldots, \succ_{s_m})$.

Note, P_i is i's strict preference relation over $S \cup \{i\}$, where $s\, P_i\, i$ means i strictly prefers s to being unassigned. Let R_i denote the "at least as good as" relation induced by P_i. Further, $\succ_s$ is a complete, irreflexive, and transitive binary priority relation over I. Thus, $i \succ_s j$ means that i has strictly higher priority at s than j.

We fix the set of agents and the set of objects throughout. The pair $(P, \succ)$ denotes a problem (or simply an economy).

The outcome of a problem is a matching, $\mu: I \to S \cup I$, where $\mu(i) \notin S \Rightarrow \mu(i) = i$ for any $i \in I$ and $|\mu^{-1}(s)| \leq 1$. We refer to $\mu(i)$ as the assignment of i under μ. A matching μ Pareto dominates matching ν, if $\mu(i)\, R_i\, \nu(i)$ for all $i \in I$ and $\mu(i)\, P_i\, \nu(i)$ for some $i \in I$. A matching is Pareto efficient if it is not Pareto dominated by any other matching.

A matching μ is blocked by an agent if it is not individually rational, meaning that there is $i \in I$ who prefers remaining unassigned to $\mu(i)$. A matching is blocked by a pair if there is an agent-object pair (i, s) where i prefers s to her assignment $\mu(i)$ and either s has not been assigned under μ or there is a lower priority agent j

who was assigned s under μ. A matching eliminates justified envy if it is not blocked by any agent or by any pair.

A mechanism φ selects a matching for each economy. Let $\varphi(P, \succ)$ denote the matching selected by φ for economy $(P, \succ)$. Let $\varphi(P, \succ)(i)$ denote the assignment of i in matching $\varphi(P, \succ)$. A mechanism is Pareto efficient if it only selects Pareto-efficient matchings. A mechanism eliminates justified envy if it only selects justified envy-free matchings. A mechanism φ is strategy-proof if reporting true preferences is a weakly dominant strategy for every agent in the preference revelation game induced by φ.

A. *Mechanisms*

Aside from Gale and Shapley's agent-proposing DA algorithm, there are two other important mechanisms for priority-based resource allocation. The first is serial dictatorship. Given a preference profile and an ordering of agents, a serial dictatorship assigns the highest ranked agent her first choice, the second highest ranked agent her top choice among remaining objects, and so on. SD is strategy-proof and Pareto efficient.

The second mechanism is Abdulkadiroğlu and Sönmez's (2003) adaptation of Gale's TTC for settings with priorities. TTC finds a matching via the following algorithm. Initially all agents and objects are available. At each step, each available agent points to her top choice among all available objects. If an agent has no acceptable objects among the remaining ones, she is assigned to herself and becomes unavailable. Each available object points to the agent who has the highest priority for the object among all available agents. There is at least one cycle. A cycle $c = \{s_k, i_k\}_{k=1,\ldots,K}$ is an ordered list of objects and agents such that s_k points to i_k and i_k points to s_{k+1} for every k, where $s_{K+1} = s_1$. Moreover, each object or agent can be part of at most one cycle. Every agent in a cycle is assigned the object she points to and is removed. The assigned object becomes unavailable. The algorithm terminates when no more agents can be assigned objects. TTC is strategy-proof and Pareto efficient.

B. *Comparing Mechanisms*

Our problem-wise comparison between mechanisms is defined as follows.

DEFINITION 1: *Mechanism φ has less justified envy than ψ at $\succ$, if for any P and agent-object pair (i, s), if pair (i, s) blocks $\varphi(P, \succ)$, then pair (i, s) blocks $\psi(P, \succ)$. A mechanism φ has less justified envy than ψ if it has less justified envy than ψ at each $\succ$. A mechanism φ has strictly less justified envy than ψ if φ has less justified envy than ψ, but ψ does not have less justified envy than φ.*

If φ has less justified envy than ψ, then the set of agent-object blocking pairs under φ is a subset of that under ψ for each problem.

DEFINITION 2: *Given a class of mechanisms $\mathcal{C}$, φ is justified envy minimal in $\mathcal{C}$ if there is no other mechanism ψ in $\mathcal{C}$ that has strictly less justified envy than φ.*

Our problem-wise comparison is related to Chen and Kesten (2017), who compare mechanisms according to problems where they produce stable outcomes. Our notion allows us to also compare mechanisms that needn't produce stable outcomes.

II. TTC as Justified Envy Minimizer

Our main result establishes the following constrained-optimality property of TTC in one-to-one matching.

THEOREM 1: *TTC is justified envy minimal in the class of Pareto-efficient and strategy-proof mechanisms.*

PROOF:

Let φ be a Pareto-efficient and strategy-proof mechanism. We show that if φ has less justified envy than *TTC* at $\succ$, then $\varphi(\cdot,\succ) = TTC(\cdot,\succ)$.

To the contrary, suppose there exists P such that

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

Let $I_k(P,\succ)$ be the set of agents who are matched in step k of $TTC(P,\succ)$, and let ℓ be the smallest k such that $I_k(P,\succ)$ contains an agent who is assigned differently between *TTC* and φ. By definition, for some $i \in I_\ell(P,\succ)$,

$$\varphi(P,\succ)(i) \neq TTC(P,\succ)(i).$$

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

Consider the preference relation:

$$P'_{i_K} : s_1, s_K, \dots$$

Since we've only altered the preferences of i_K in c and $TTC(P,\succ)(i_K) = s_1$ has become her first choice, the *TTC* matching remains the same:

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

Since

$$\varphi(P,\succ)(i_K) \neq \underbrace{TTC(P,\succ)(i_K)}_{=s_1},$$

and $\varphi(P,\succ)(i_K)$ is still available at step ℓ, we obtain

$$s_1 \; P_{i_K} \; \varphi(P,\succ)(i_K). \tag{1}$$

But since $TTC\left(P'_{i_K}, P_{-\{i_K\}}, \succ\right)(i_K) = s_1 \; P'_{i_K} \; s_K$, (i_K, s_K) does not block $TTC\left(P'_{i_K}, P_{-\{i_K\}}, \succ\right)$.

Since φ has less justified envy than TTC at $\succ$, (i_K, s_K) should not block $\varphi(P'_{i_K}, P_{-\{i_K\}}, \succ)$. Since s_K points to i_K in cycle c, $i_K \succ_{s_K} j$ for all $j \neq i_K$ who are still unassigned at step ℓ. Given that s_K will be assigned one of the agents still unassigned at step ℓ, by construction, we must then have

$$\varphi\big(P'_{i_K}, P_{-\{i_K\}}, \succ\big)(i_K) \in \{s_1, s_K\}.$$

Given (1), strategy-proofness of φ implies

$$\varphi\big(P'_{i_K}, P_{-\{i_K\}}, \succ\big)(i_K) = s_K;$$

otherwise i_K would be able to manipulate φ in economy $(P, \succ)$ by submitting P'_{i_K} to obtain s_1. Hence, i_K obtains her second choice under P'_{i_K}. Now we have $\varphi\big(P'_{i_K}, P_{-\{i_K\}}, \succ\big)(i_{K-1}) \neq s_K = TTC\big(P'_{i_K}, P_{-\{i_K\}}, \succ\big)(i_{K-1})$, and so we can apply the same argument for i_{K-1}.

By iterating the argument for $i_{K-1}, \ldots, i_1$, and for every agent in the cycle c, we obtain that

$$\varphi(P'_c, P_{-c}, \succ)(i_k) = s_k,$$

where $P'_c = \{P'_{i_k}\}_{i_k \in c}$ and $P'_{i_k} : s_{k+1}, s_k, \ldots$.

To see this, note that at the last step of the argument, we have

$$\varphi(P'_c, P_{-c}, \succ)(i_1) = s_1 \neq \varphi(P'_c, P_{-c}, \succ)(i_K).$$

Since

$$TTC\big(P'_c, P_{-c}, \succ\big)(i_K)\ R'_{i_K}\ s_K,$$

we know (i_K, s_K) does not block $TTC\big(P'_c, P_{-c}, \succ\big)$, and because φ has less justified envy than TTC, it does not block $\varphi(P'_c, P_{-c}, \succ)$. Since, by definition of TTC, $i_K \succ_{s_K} j$, for all $j \neq i$ remaining in step ℓ of TTC, and, by construction, s_K gets matched one of the agents remaining at step ℓ,

$$\varphi(P'_c, P_{-c}, \succ)(i_K)\ R'_{i_K}\ s_K,$$

which implies

$$\varphi(P'_c, P_{-c}, \succ)(i_K) = s_K.$$

Hence, starting from $\varphi(P'_c, P_{-c}, \succ)(i_1) = s_1$, we obtain $\varphi(P'_c, P_{-c}, \succ)(i_K) = s_K$. Applying the argument iteratively, we find $\varphi(P'_c, P_{-c}, \succ)(i_k) = s_k$ for all k. But this contradicts Pareto efficiency of φ because every agent in the cycle will be better off if every i_k is matched with s_{k+1} (modulo k) without changing the matching of other agents, establishing the claim. ∎

Theorem 1 does not imply that TTC is the only justified envy–minimal mechanism among Pareto-efficient and strategy-proof mechanisms. However, among a subclass that includes SD, TTC is the only justified envy–minimal mechanism among Pareto-efficient and strategy-proof mechanisms. To define that subclass, for each object s, let $f_s(\succ)$ be a priority relation for s obtained by a function of $\succ$, and $f = (f_s)_s$. The subclass consists of priority-adjusted TTC mechanisms, $\varphi(\cdot,\cdot) = TTC(\cdot, f(\cdot))$, which are obtained by running TTC under priorities $f(\cdot)$ that may differ from true priorities. Observe that every φ in this class is Pareto efficient and strategy-proof. Serial dictatorship is obtained if each object's priority is the same.

PROPOSITION 1: *Suppose* $f_s(\succ) \neq \succ_s$ *for some* s. *Then, the mechanism* $\varphi(\cdot,\cdot) = TTC(\cdot, f(\cdot))$ *is not justified envy minimal in the class of Pareto-efficient and strategy-proof mechanisms.*

PROOF:

Assume there is an object $s \in S$ for which $\succ'_s := f_s(\succ_s) \neq \succ_s$. This means there are at least two agents i_1 and i_2 such that $i_1 \succ_s i_2$, but $i_2 \succ'_s i_1$. Let $\succ_s$ be the priorities of object s, and further assume all objects s' have the same priorities as object s. This means DA is Pareto efficient at this profile of priorities. Clearly, DA is a strategy-proof, Pareto-efficient, and justified envy-free mechanism at this profile of priorities. Now, consider a profile of preferences where all agents except i_1 and i_2 rank all objects as unacceptable. In turn, i_1 and i_2 find only object s acceptable. Obviously in such a case, the "distorted" TTC fails to select a justified envy-free assignment. Hence, $DA(\cdot,\succ)$ has strictly less justified envy than $\varphi(\cdot,\succ)$ at $\succ$. This implies the "distorted" TTC cannot be justified envy minimal. ∎

This result shows that the justified envy minimality of TTC is nontrivial. It also highlights the roles played by the priorities: justified envy minimality fails with a minimal modification of priorities. This result clarifies the role of priorities in TTC relative to an important class of other Pareto-efficient and strategy-proof mechanisms.

Our characterization is also "tight" in the sense that relaxing any of the assumptions yields a mechanism with less justified envy and the remaining properties. When we drop efficiency, TTC does not minimize justified envy among strategy-proof mechanisms. DA does.

When we drop strategy-proofness, TTC also does not minimize envy in the class of Pareto-efficient mechanisms. Consider a mechanism φ that produces $DA(P,\succ)$ when $DA(P,\succ)$ is Pareto efficient and produces $TTC(P,\succ)$ otherwise. Clearly, φ is Pareto efficient and has less justified envy than TTC. Our main result implies φ is not strategy-proof.[3]

III. Many-to-One Matching

Models of one-to-one matching account for important applications such as housing and organs. Other applications, notably school assignment, involve many-to-one matching. Our model can be extended to this setting by assuming that each s_i

[3] Kesten (2010) shows that there is no mechanism that is Pareto efficient and strategy-proof and selects a matching eliminating justified envy whenever it exists.

has $q_{s_i} \in \mathbb{N}$ copies and that a matching must satisfy $|\mu^{-1}(s_i)| \leq q_{s_i}$ for each $s_i \in S$. TTC is also easily adapted to the many-to-one environment with a simple modification (see Abdulkadiroğlu and Sönmez 2003): when an object is in a cycle, a single *copy* is assigned, and the object becomes unavailable when all its copies are assigned. The resulting mechanism continues to be efficient and strategy-proof.

The following example illustrates that TTC is not justified envy minimal in this setting.

EXAMPLE 1: *There are three agents* $I = \{i_1, i_2, i_3\}$ *and two objects* $S = \{s_1, s_2\}$, *with* $q_{s_1} = 2$ *and* $q_{s_2} = 1$. *The preferences and priorities are given by*

P_{i_1}	P_{i_2}	P_{i_3}	$\succ_{s_1}$	$\succ_{s_2}$
s_2	s_1	s_2	i_1	i_2
s_1	s_2	s_1	i_2	i_3
			i_3	i_1

TTC produces

$$\begin{pmatrix} i_1 & i_2 & i_3 \\ s_2 & s_1 & s_1 \end{pmatrix},$$

where agent i_3 *has justified envy and is unable to get* s_2 *despite having a higher priority than* i_1. *If DA were used for the above priority profile, it would eliminate justified envy but also produce efficient matchings, regardless of the individuals' preferences.*[4] *If one were to use DA only for the above priority and TTC for the other priorities, then the resulting mechanism is Pareto efficient, strategy-proof, and has strictly less justified envy than TTC.*

With multiple copies available for a given object, multiple agents effectively have top priority at that object; in the above example, both i_1 and i_2 are guaranteed to be assigned s_1. Unlike one-to-one matching, an agent can be associated with multiple cycles, and which cycle is cleared first matters. For instance, if i_2 were assigned through a self-cycle instead of one with i_1, the second matching would have been obtained. Unfortunately, no general method is known for clearing the "right" cycles, and thus there is no obvious practical mechanism that is efficient, strategy-proof and justified envy minimal.[5]

In light of this, one may ask whether TTC still reduces justified envy in many-to-one matching in comparison with SD. It is not possible to compare the two mechanisms uniformly across all priorities and all serial orders for a given preference profile. Therefore, we compare mechanisms in terms of "average" incidences

[4] This follows since the priority structure is acyclic in the sense of Ergin (2002), a sufficient condition for DA to be Pareto efficient. At the same time, the above priority structure is not acyclic in the sense of Kesten (2006), as implied by justified envy present in TTC. In this sense, this example exploits the "gap" between the Ergin acyclicity and Kesten acyclicity. Although these conditions are suggestive of why many-to-one matching differs from one-to-one matching, they do not explain why TTC is justified envy minimal, as claimed by Theorem 1, in one-to-one matching, since that result applies even when Ergin acyclicity fails.

[5] Morrill (2015) proposes a variant of TTC where self-cycles are cleared first. This variant as well as others proposed in the literature, such as the equitable top trading cycle mechanism proposed by Hakimov and Kesten (2014), fail to be justified envy minimal, however, as discussed in our working paper.

of justified envy, where the average is taken with respect to all possible priorities and all possible serial orders under which SD may be run for a given preference profile. It is analytically more convenient to take a "probabilistic" perspective in which TTC is run with the priorities randomly drawn according to the uniform distribution and, given a priority profile, SD is run with its serial order drawn according to the uniform distribution. Note that the latter mechanism corresponds to the random serial dictatorship (RSD). For comparison in the "average" sense, we consider whether one mechanism dominates another in a probabilistic sense.

Consider the many-to-one matching model in which objects' priorities $\succ$ are randomly generated from draws of a uniform distribution. Further assume that the total number of objects' copies equals the number of individuals (i.e., $\sum_k q_{s_k} = n$).[6] To state our next result, we introduce notation that counts blocking pairs for a given agent i and object s. For mechanism $\phi \in \{TTC, RSD\}$ and any pair (i,s) where i is assigned to an object ranked lower than s, let $N^{\phi}(i,s)$ denote the number of agents assigned object s with lower priority at s than i. Since priorities are random, $N^{\phi}(i,s)$ is a random variable.

We compare $N^{\phi}(i,s)$ between TTC and RSD as follows.

THEOREM 2: *Given an agent-object pair (i,s), $N^{RSD}(i,s)$ first-order stochastically dominates $N^{TTC}(i,s)$, that is, for any $\ell \geq 0$,*

$$\Pr\{N^{TTC}(i,s) \geq \ell\} \leq \Pr\{N^{RSD}(i,s) \geq \ell\}. \tag{2}$$[7]

The proof is in the Appendix. This theorem shows that the number of priorities under which i justifiably envies ℓ or more agents over s in the TTC assignment is weakly less than the expected number of priorities under which i justifiably envies ℓ or more agents over s in the RSD assignment. The intuition boils down to the following observation. Unlike RSD, which ignores priorities, TTC uses priorities. But does it use them to reduce justified envy? The answer depends on the types of cycles through which an agent is assigned. If an agent is assigned via a short cycle—that is, the agent points to an object that in turn points back to that agent—then no justified envy occurs since any agent who may envy her will still remain at the round of assignment and, hence, must have a lower priority than her. Meanwhile, if an agent is assigned via a longer cycle, one that involves more than one agent, then the agent assigned that object has the same probability of having higher priority than those who may envy her. TTC assigns some agents via short cycles, which do not create justified envy. Agents assigned via long cycles create the same amount of justified envy as RSD. Theorem 2 immediately allows us to establish the following corollary.

[6] In this context, Pathak and Sethuraman (2011) show that agents have the same random allocation under TTC with random priorities and under random serial dictatorship. This equivalence does not imply that both mechanisms are equivalent in terms of justified envy. To see this, suppose two agents rank two objects the same. Under TTC, whichever object they both prefer is assigned to the agent with the highest priority at that object. Hence, no justified envy ever arises. But under random serial dictatorship, the commonly preferred object may be assigned to the agent with lower priority with positive probability, in which case the higher priority agent would have justified envy.

[7] In addition, if, under TTC, agent i prefers object s to his assignment with strictly positive probability, then the inequality is strict for each $\ell = 1, \ldots, q_s$.

COROLLARY 1: *When priorities are drawn uniformly randomly, the expected number of agents with justified envy, the expected number of blocking pairs, and the expected number of agents each agent justifiably envies are all smaller under TTC than under RSD.*

This result formalizes how TTC admits less justified envy than serial dictatorship "on average" when objects may have multiple copies. In our working paper, we show our result extends when we relax the assumption of uniformly randomly drawn priorities to accommodate priorities that are correlated across objects and location-based priorities. Moreover, this result persists as the economy grows large, when the number of copies per object are large relative to the number of objects as in Abdulkadiroğlu, Che, and Yasuda (2015) and Azevedo and Leshno (2016). Under this asymptotic sequence, using our characterization results in the Appendix (Propositions 2 and 3), we can show that the expected number of agents each agent justifiably envies is strictly smaller under TTC than under RSD in the limit economy when the probability measure over preferences has full support.[8]

IV. Comparing Mechanisms in New Orleans and Boston

While Corollary 1 provides a clear comparison of TTC versus RSD in an average sense for all possible priorities, in reality priorities are determined by applicants' home address, siblings' enrollment status, and other criteria. We therefore compare TTC and RSD with actual priorities using data from New Orleans and Boston (New Orleans Recovery School District 2012b, Boston Public Schools 2013).

In 2011–2012, the New Orleans Recovery School District pioneered a unified enrollment process called OneApp, integrating admissions to all types of schools under a single offer system. Officials identified three major priority groups: sibling, applying from a closing school, and geography. The discussion about mechanism centered on the trade-off between efficiency and eliminating justified envy, and eventually TTC was selected based on the desire for "as many students as possible to get into their top choice school" (New Orleans Recovery School District 2012a). Vanacore (2011) and Vanacore (2012) provide additional details.

We use data from elementary and high school applicants in grades prekindergarten (PK) and nine in the 2011–2012 school year and compare TTC, RSD, and DA. There are a total of 46 schools assigned to students in grades PK and nine. Applicants can rank up to eight choices; in practice, less than 5 percent of applicants rank 8 choices, so we simulate the algorithms holding preferences fixed since truth telling is a weakly dominant strategy in each mechanism without a constraint.[9]

[8] However, when the number of copies is fixed and the number of individuals and objects grow at the same rate, TTC may not admit significantly less justified envy than RSD. More specifically, Che and Tercieux (2018) study TTC in a one-to-one matching model in which both applicants' preferences and their priorities are randomly drawn according to the uniform distribution. They show that, as the economy grows large, the incidences of justified envy under TTC and RSD become indistinguishable. A rough intuition is that as the economy grows arbitrarily large, most of the agents are assigned via long cycles at least in the one-to-one matching, so the distinction between TTC and RSD vanishes. Under these assumptions, the advantage of TTC over RSD vanishes asymptotically.

[9] Strictly speaking, when the number of choices is constrained, none of the mechanisms is strategy-proof. But for the 95 percent of applicants who ranked fewer than eight choices, they could have ranked an additional

Table 1—Comparison of Mechanisms in New Orleans and Boston

	New Orleans OneApp			Boston		
	TTC	SD	Student-proposing DA	TTC	SD	Student-proposing DA
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. Choice assigned						
1	772	777	762	1,240	1,236	1,227
2	126	121	137	322	315	336
3	46	44	51	134	132	138
4	18	17	19	56	51	57
5+	11	8	10	39	34	40
Unassigned	222	228	217	102	124	96
Total	1,196	1,196	1,196	1,893	1,893	1,893
Panel B. Statistics on blocking pairs						
Students with justified envy	158	213	0	129	280	0
Schools involved in blocking pairs	7	12	0	18	44	0
Blocking pairs (i,s)	228	308	0	160	369	0
Instances of justified envy $(i,(j,s))$	1,111	6,546	0	768	3,650	0

Notes: In New Orleans, the data are from 2011–2012 and grades PK and nine. In Boston, the data cover four school years from 2009–2010 through 2012–2013 and grades K1, K2, six, and nine. TTC is defined in Abdulkadiroğlu and Sonmez (2003). A student has justified envy if there exists a school s where (i,s) is a blocking pair. School s involved in blocking pairs means there is a school s such that there exists student i such that (i,s) is a blocking pair. Blocking pairs (i,s) means there exists at least one applicant j such that $(i,(j,s))$ is a blocking instance. Instance of justified envy $(i,(j,s))$ means student i has justified envy for student j's assignment at s. The numbers represent averages of 100 different lottery draws for each grade. In New Orleans, there are a total of 46 schools in PK and grade nine. The standard deviation across lottery draws in column 1 for first choice assigned is 5.2, for unassigned is 5.2, for students with justified envy is 8.5, for schools involved in blocking pairs is 0.4, for blocking pairs is 20.2, and for instances of justified envy is 110.7. Standard deviations are similar for columns 2 and 3. In Boston, there is an average of 124 schools across the four years. The standard deviation across lottery draws in column 4 for first choice assigned is 6.8, for unassigned is 5.2, for students with justified envy is 16.9, for schools involved in blocking pairs is 2.4, for blocking pairs is 21.6, and for instances of justified envy is 98.6. Standard deviations are similar for columns 5 and 6.

We fix preferences and use the actual priorities. The priorities are coarse, and lotteries are used to break ties in priorities. We draw 100 sets of lottery numbers, one for each applicant, and run the assignment algorithms for each lottery draw. For SD, we order applicants according to the realized lottery order. Table 1 reports the average across lottery draws and two grades. For comparison purposes, we also report the corresponding numbers in Boston Public Schools.

Both TTC and SD assign more applicants to their first choices than DA, but the difference in the aggregate rank distribution is small.[10] This comes at the cost of creating instances of justified envy. Whereas DA is justified envy free, on average in New Orleans, 213 students exhibit justified envy under SD in comparison to 158 students under TTC. The number of schools involved in a blocking pair increases from 7 under TTC to 12 under SD. Strikingly, the number of justified envy instances at which a student may make a complaint against another student at a more preferred

choice, and so there is little reason to strategize among the choices they submitted (see Haeringer and Klijn 2009 and Pathak and Sönmez 2013).

[10] Unlike New Orleans and Boston, there is a substantial difference in the aggregate rank distribution between TTC and DA in NYC (see Abdulkadiroğlu, Pathak, and Roth 2009 and Che and Tercieux 2019).

school is almost six times as high under SD as under TTC. Our empirical analysis in Boston yields similar results.

This evidence complements Theorem 1 and Corollary 1 in supporting the idea that TTC performs well in terms of "economizing" on justified envy when schools have more than one seat and under real-world priority structures.

V. Conclusion

This paper provides arguments for TTC over other Pareto-efficient and strategy-proof mechanisms for priority-based resource allocation. Our main result is a counterpart to DA's constrained-optimality result for the one-to-one matching model. We also show that TTC outperforms SD in an average sense for the many-to-one environment. Finally, using data from New Orleans and Boston, we also show that TTC has significantly less justified envy than SD with real-world priorities and preferences.

In the field, there is growing momentum for DA over TTC (see Abdulkadiroğlu 2013 and Pathak 2017). This trend may be driven by a first-mover advantage of DA and its use in other contexts. New York City and Boston adopted DA in 2003 and 2005, and DA is widely used in residency matching (Roth and Peranson 1999). In 2013, New Orleans also switched from TTC to DA. One of the most important reasons for this switch involved challenges in explaining how TTC handles priorities. Under DA, officials could explain that an applicant did not obtain an assignment at a higher ranked seat because another applicant with higher priority was assigned to that seat. At the time of the change, a clear explanation of how TTC reflects priorities was not available.

It remains to be seen whether TTC will be used in the field again. But policymakers cannot ignore efficiency, which TTC delivers but DA does not. For this reason, TTC should remain a serious policy option. Our formal results may make it easier to explain how TTC incorporates priorities. It's possible that TTC would have been chosen in some settings with knowledge of this result, and at the very least, advocates now have a new argument in its favor.

Appendix: Proof of Theorem 2

Given an agent i, let $TTC(i)$ denote the assignment of agent i. Given an object s, we abuse notation and let $TTC(s)$ denote the set of agents assigned to object s. Recall that a *short cycle* is a cycle in which an agent points to an object and the object points to that agent at some step of TTC. A *long cycle* is a cycle that is not short.

For some agent $j \in TTC(s)$, let $k(j)$ be the agent pointed to by object s when agent j is part of a cycle (if this is a short cycle, then $k(j) = j$). Let $\underline{k}$ denote the agent in $\bigcup_{j \in TTC(s)} k(j)$ with the lowest priority.

The set of agents in $TTC(s)$ can be partitioned into two sets $J_1 \cup J_2$:

J_1: agents in J_1 have a weakly higher priority than $\underline{k}$, and

J_2: agents in J_2 have a strictly lower priority than $\underline{k}$.

Here, $N^{TTC}(i,s)$ is the number of agents receiving object s that agent i justifiably envies under TTC. We consider the event $\{N^{TTC}(i,s) \geq \ell\}$ that agent i justifiably envies ℓ or more agents assigned to object s by TTC. Also let $\mathcal{E}_i^{TTC}(s)$ denote the event that agent i prefers s to $TTC(i)$.

PROPOSITION 2 (A TTC Characterization Result): *Given any* (i,s) *and* $\ell \geq 1$,

$$\Pr\left\{N^{TTC}(i,s) \geq \ell \,\middle|\, \mathcal{E}_i^{TTC}(s)\right\} = E\left[\max\left\{1 - \frac{\ell}{|J_2|+1}, 0\right\} \,\middle|\, \mathcal{E}_i^{TTC}(s)\right].$$

In addition,

$$\Pr\left\{N^{TTC}(i,s) \geq \ell \,\middle|\, \neg\mathcal{E}_i^{TTC}(s)\right\} = 0.$$

PROOF:

If agent i does not prefer s to $TTC(i)$ (that is, $\mathcal{E}_i^{TTC}(s)$ does not hold), she never justifiably envies an agent assigned s. If $\mathcal{E}_i^{TTC}(s)$ holds, there may be positive probability that agent i has justified envy toward an agent assigned s. The following claim is the main step of the proof.

CLAIM 1: *Fix* $\ell \geq 1$. *For any* $J \subset I$, *we must have*

$$\Pr\left\{N^{TTC}(i,s) \geq \ell \,\middle|\, \mathcal{E}_i^{TTC}(s) \text{ and } J_2 = J\right\} = \max\left\{1 - \frac{\ell}{|J|+1}, 0\right\}.$$

PROOF:

Let $J \equiv \{j_1, \ldots, j_{|J|}\}$. We first show that

$$\begin{aligned} (3) \qquad & \Pr\left\{J_2 = \{j_1, \ldots, j_{|J|}\}, i \succ_s j_1 \succ_s \cdots \succ_s j_{|J|}, \mathcal{E}_i^{TTC}(s)\right\} \\ & = \Pr\left\{J_2 = \{j_1, \ldots, j_{|J|}\}, \pi(i) \succ_s \pi(j_1) \succ_s \cdots \succ_s \pi(j_{|J|}), \mathcal{E}_i^{TTC}(s)\right\}, \end{aligned}$$

where $\pi : J \cup \{i\} \to J \cup \{i\}$ is an arbitrary permutation of $J \cup \{i\}$.

Consider any realization of priorities $\succ$ under which

$$\left\{J_2 = \{j_1, \ldots, j_{|J|}\}, i \succ_s j_1 \succ_s \cdots \succ_s j_{|J|}, \mathcal{E}_i^{TTC}(s)\right\}$$

holds. Note first that any agent in J_2 must be assigned via a long cycle. If some agent $j \in J_2$ is assigned via a short cycle, then $j = k(j)$. Since by definition of $\underline{k}$, $k(j) \succeq_s \underline{k}$, we have $j \succeq_s \underline{k}$. But then agent j cannot belong to J_2 by construction of J_2.

Next, recall that, by definition, $\underline{k} \succ_s j$ for any $j \in J_2$. Moreover, $\underline{k} \succ_s i$ since $\mathcal{E}_i^{TTC}(s)$ holds. Hence, if we start from $\succ$ and permute the priority of agent i and agents in J_2 according to π, then TTC yields exactly the same assignment as under the permuted priority. This is because at any step of TTC involving object s before the permutation, the cycle that is formed involving s is unchanged after the permutation. Since object s points to $k(j) \succeq_s \underline{k}$ before the permutation and since $\underline{k} \succ_s i$ and $\underline{k} \succ_s j$ for any $j \in J$, s continues to point to $k(j)$ after the permutation.

After permutation π, the set J_2 consists of exactly the same agents since the permutation leaves all the steps of TTC unchanged and since this permutation is restricted to agents in $J \cup \{i\}$ who all have a strictly lower priority than $\underline{k}$, which leaves $\underline{k}$ unchanged after the permutation. Hence, after the permutation, the event

$$\left\{J_2 = \{j_1, \ldots, j_{|J|}\}, \pi(i) \succ_s \pi(j_1) \succ_s \cdots \succ_s \pi(j_{|J|}), \mathcal{E}_i^{TTC}(s)\right\}$$

occurs, and the sets J_1 and J_2 must also remain the same.

Let $\mathcal{P}_1$ and $\mathcal{P}_2$ be the set of priorities giving rise to

$$\left\{J_2 = \{j_1, \ldots, j_{|J|}\}, i \succ_s j_1 \succ_s \cdots \succ_s j_{|J|}, \mathcal{E}_i^{TTC}(s)\right\}$$

and

$$\left\{J_2 = \{j_1, \ldots, j_{|J|}\}, \pi(i) \succ_s \pi(j_1) \succ_s \cdots \succ_s \pi(j_{|J|}), \mathcal{E}_i^{TTC}(s)\right\},$$

respectively. We have constructed an injection from $\mathcal{P}_1$ to $\mathcal{P}_2$ and so $|\mathcal{P}_1| \leq |\mathcal{P}_2|$. This follows from the fact that distinct priority profiles within $\mathcal{P}_1$ yield distinct priority profiles in $\mathcal{P}_2$ upon permuting.

Using a similar argument, we can also build an injection from $\mathcal{P}_2$ to $\mathcal{P}_1$ to show that

$$|\mathcal{P}_1| = |\mathcal{P}_2|.$$

Since objects' priorities $\succ$ are uniform random, (3) follows.

Consequently, conditional on $\mathcal{E}_i^{TTC}(s)$ and $\{J_2 = J\}$, agent i has a lower priority than any agent in $TTC(s) \backslash J$, and the priority ordering of agents $\{i\} \cup J$ by s is uniform random. Hence, for $\ell \leq |J|$, the conditional probability that i justifiably envies ℓ or more agents obtaining s, or equivalently the conditional probability that i is not among the ℓth lowest priority agents for s, is

$$1 - \frac{\ell}{|J| + 1}.$$

Obviously, the conditional probability is 0 for $\ell > |J|$. ∎

Using the claim, we have

$$\begin{aligned}
&\Pr\left\{N^{TTC}(i,s) \geq \ell \middle| \mathcal{E}_i^{TTC}(s)\right\} \\
&\quad = \sum_J \Pr\left(J_2 = J \middle| \mathcal{E}_i^{TTC}(s)\right) \Pr\left\{N^{TTC}(i,s) \geq \ell \middle| \mathcal{E}_i^{TTC}(s), J_2 = J\right\} \\
&\quad = \sum_J \Pr\left(J_2 = J \middle| \mathcal{E}_i^{TTC}(s)\right) \max\left\{1 - \frac{\ell}{|J| + 1}, 0\right\} \\
&\quad = E\left[\max\left\{1 - \frac{\ell}{|J_2| + 1}, 0\right\} \middle| \mathcal{E}_i^{TTC}(s)\right],
\end{aligned}$$

which completes the proof. ∎

We also have an analogous result for RSD. Let $\mathcal{E}_i^{RSD}(s)$ be the event that agent i prefers s to $RSD(i)$, and recall that $N^{RSD}(i,s)$ stands for the number of agents assigned to object s that agent i justifiably envies under RSD.

PROPOSITION 3: *Given any* (i,s) *and* $\ell \geq 1$, *we have*

$$\Pr\left\{N^{RSD}(i,s) \geq \ell \,\middle|\, \mathcal{E}_i^{RSD}(s)\right\} = \max\left\{1 - \frac{\ell}{q_s + 1}, 0\right\}.$$

In addition,

$$\Pr\left\{N^{RSD}(i,s) \geq \ell \,\middle|\, \neg\mathcal{E}_i^{RSD}(s)\right\} = 0.$$

Propositions 2 and 3 allow us to prove Theorem 2.

PROOF OF THEOREM 2:
Fix any positive integer $\ell \geq 1$ (the argument is trivial for $\ell = 0$). Then

$$\begin{aligned}
\Pr\left\{N^{TTC}(i,s) \geq \ell\right\} &= \Pr\left\{\mathcal{E}_i^{TTC}(s)\right\}\Pr\left\{N^{TTC}(i,s) \geq \ell \,\middle|\, \mathcal{E}_i^{TTC}(s)\right\} \\
&= \Pr\left\{\mathcal{E}_i^{TTC}(s)\right\} E\left[\max\left\{1 - \frac{\ell}{|J_2| + 1}, 0\right\} \,\middle|\, \mathcal{E}_i^{TTC}(s)\right] \\
&\leq \Pr\left\{\mathcal{E}_i^{TTC}(s)\right\}\max\left\{1 - \frac{\ell}{q_s + 1}, 0\right\} \\
&= \Pr\left\{\mathcal{E}_i^{RSD}(s)\right\}\max\left\{1 - \frac{\ell}{q_s + 1}, 0\right\} \\
&= \Pr\left\{\mathcal{E}_i^{RSD}(s)\right\}\Pr\left\{N^{RSD}(i,s) \geq \ell \,\middle|\, \mathcal{E}_i^{TTC}(s)\right\} \\
&= \Pr\left\{N^{RSD}(i,s) \geq \ell\right\},
\end{aligned}$$

where the first and second equalities use Proposition 2, the first inequality uses the fact that $|J_2| \leq q_s$, the third equality uses the Pathak and Sethuraman (2011) equivalence result, while the last two equalities use Proposition 3. ∎

REFERENCES

Abdulkadiroğlu, Atila. 2013. "School Choice." In *The Handbook of Market Design*, edited by Nir Vulkan, Alvin E. Roth, and Zvika Neeman, chapter 5. Oxford: Oxford University Press.

Abdulkadiroğlu, Atila, Yeon-Koo Che, and Yosuke Yasuda. 2015. "Expanding 'Choice' in School Choice." *American Economic Journal: Microeconomics* 7 (1): 1–42.

Abdulkadiroğlu, Atila, Parag A. Pathak, and Alvin E. Roth. 2009. "Strategy-Proofness versus Efficiency in Matching with Indifferences: Redesigning the NYC High School Match." *American Economic Review* 99 (5): 1954–78.

Abdulkadiroğlu, Atila, Yeon-Koo Che, Parag A. Pathak, Alvin E. Roth, and Olivier Tercieux. 2020. "Efficiency, Justified Envy, and Incentives in Priority-Based Matching: Dataset." *American Economic Review: Insights*. https://doi.org/10.1257/aeri.20190307.

Abdulkadiroğlu, Atila, and Tayfun Sönmez. 2003. "School Choice: A Mechanism Design Approach." *American Economic Review* 93 (3): 729–47.

Abdulkadiroğlu, Atila, Parag A. Pathak, Alvin E. Roth, and Tayfun Sönmez. 2005. "The Boston Public School Match." *American Economic Review* 95 (2): 368–71.

Azevedo, Eduardo M., and Jacob D. Leshno. 2016. "A Supply and Demand Framework for Two-Sided Matching Markets." *Journal of Political Economy* 124 (5): 1235–68.

Boston Public Schools. 2013. "Student Administrative Data System." Boston: Boston Public Schools.

BPS Strategic Planning Team. 2005. "Recommendation to Implement a New BPS Assignment Algorithm." Presentation to the Boston School Committee, Boston.

Che, Yeon-Koo, and Olivier Tercieux. 2018. "Top Trading Cycles in Prioritized Matching: An Irrelevance of Priorities in Large Markets." https://sites.google.com/site/otercieux/ (accessed September 19, 2020).

Che, Yeon-Koo, and Olivier Tercieux. 2019. "Efficiency and Stability in Large Matching Markets." *Journal of Political Economy* 127 (5): 2301–41.

Chen, Yan, and Onur Kesten. 2017. "Chinese College Admissions and School Choice Reforms: A Theoretical Analysis." *Journal of Political Economy* 125 (1): 99–139.

Dubins, L. E., and D. A. Freedman. 1981. "Machiavelli and the Gale-Shapley Algorithm." *American Mathematical Monthly* 88 (7): 485–94.

Dur, Umut Mert. 2013. "A Characterization of the Top Trading Cycles Mechanism for the School Choice Problem." https://sites.google.com/site/umutdur/research.

Ergin, Haluk I. 2002. "Efficient Resource Allocation on the Basis of Priorities." *Econometrica* 70 (6): 2489–98.

Gale, David, and Lloyd S. Shapley. 1962. "College Admissions and the Stability of Marriage." *American Mathematical Monthly* 69 (1): 9–15.

Haeringer, Guillaume, and Flip Klijn. 2009. "Constrained School Choice." *Journal of Economic Theory* 144 (5): 1921–47.

Hakimov, Rustamdjan, and Onur Kesten. 2014. "The Equitable Top Trading Cycles Mechanism for School Choice." Wissenschaftszentrum Berlin für Sozialforschung Discussion Paper SP II 2014-210.

Kesten, Onur. 2006. "On Two Competing Mechanisms for Priority-Based Allocation Problems." *Journal of Economic Theory* 127 (1): 155–71.

Kesten, Onur. 2010. "School Choice with Consent." *Quarterly Journal of Economics* 125 (3): 1297–1348.

Landsmark, Ted, Helen Dajer, and Lisa Gonsalves. 2004. *Report and Recommendations of the Boston Public Schools Student Assignment Task Force*. Boston: Boston School Committee.

Ma, Jinpeng. 1994. "Strategy-Proofness and the Strict Core in a Market with Indivisibilities." *International Journal of Game Theory* 23: 75–83.

Morrill, Thayer. 2015. "Making Just School Assignments." *Games and Economic Behavior* 92: 18–27.

New Orleans Recovery School District. 2012a. "Overview of the New Public School Enrollment Process." New Orleans Recovery School District Presentation to Public.

New Orleans Recovery School District. 2012b. "Student Administrative Data System."

Pápai, Szilvia. 2000. "Strategyproof Assignment by Hierarchical Exchange." *Econometrica* 68 (6): 1403–33.

Pathak, Parag A. 2017. "What Really Matters in Designing School Choice Mechanisms." In *Advances in Economics and Econometrics: Eleventh World Congress of the Econometric Society*, edited by Larry Samuelson, 176–214. Cambridge, UK: Cambridge University Press.

Pathak, Parag A., and Jay Sethuraman, 2011. "Lotteries in Student Assignment: An Equivalence Result." *Theoretical Economics* 6 (1): 1–17.

Pathak, Parag A., and Tayfun Sönmez. 2013. "School Admissions Reform in Chicago and England: Comparing Mechanisms by Their Vulnerability to Manipulation." *American Economic Review* 103 (1): 80–106.

Pycia, Marek, and Ünver. 2017. "Incentive Compatible Allocation and Exchange of Discrete Resources." *Theoretical Economics* 12 (1): 287–329.

Roth, Alvin E. 1982a. "Incentive Compatibility in a Market with Indivisible Goods." *Economics Letters* 9 (2): 127–32.

Roth, Alvin E. 1982b. "The Economics of Matching: Stability and Incentives." *Mathematics of Operations Research* 7 (4): 617–28.

Roth, Alvin E., and Elliott Peranson. 1999. "The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design." *American Economic Review* 89 (4): 748–80.

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

Vanacore, Andrew. 2011. "Recovery School District Chief Plans Central Enrollment System, Technical Training, More." *New Orleans Times–Picayune*, September 7. https://www.nola.com/news/education/article_8a072c41-b2ce-5608-8a39-b33245850b80.html.

Vanacore, Andrew. 2012. "Centralized Enrollment in Recovery School District Gets First Tryout." *New Orleans Times–Picayune*, April 16. https://www.nola.com/news/education/article_e4575979-70e0-5b81-ad7f-b87ef293795c.html.

This article has been cited by:

1. Philip J. Reny. 2022. Efficient Matching in the School Choice Problem. *American Economic Review* **112**:6, 2025-2043. [Abstract] [View PDF article] [PDF with links]